

container **d**

# containerd update

**FOSDEM 2019 - Containers Devroom**

Sunday, 3 February 2019

Phil Estes ([IBM Cloud, Distinguished Engineer](#))

# Containerd Overview



'...but many platform builders and operators are looking for "[boring infrastructure](#)": a basic component that provides the robust primitives for running containers on their system, bundled in a stable interface, and nothing else.' - Docker Blog, containerd announcement

- **Project Goal:** Provide a “*{boring} core container runtime*” for the industry, allowing Docker and Kubernetes to both have a strong base on which to innovate at higher layers

## KEY TENETS

- Reliability/Stability

- Clean/usable client API

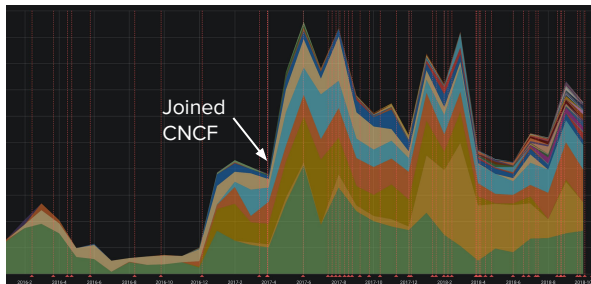
- Strong Compatibility Guarantees

- Performance



# Containerd Community

- **3,418 GH stars** (up 1000+)
- **165 contributors** (up 50+)
- **12 maintainers**, representing 8 organizations/vendors
- **6 reviewers**, representing 5 organizations/vendors
- **720** Slack channel members
- **3,441** Twitter followers (@containerd)



A screenshot of the GitHub repository page for containerd. The page shows the repository name 'containerd / containerd' and various statistics: 94 issues, 20 pull requests, 0 ZenHub projects, and 0 insights. The repository is described as 'An open and reliable container runtime' with a link to 'https://containerd.io'. It has 4,135 commits, 5 branches, 53 releases, 1 environment, 154 contributors, and is licensed under Apache-2.0. The page also shows a list of recent pull requests, including a merge pull request #2754 from mikebrow/readme-slack-invite-update.

Branch	Commit	Message	Time
master	b043f9f	Latest commit	3 hours ago
.github		Add github issue template	10 months ago
api		Prepare 1.2 release	7 days ago
archive		archive: add support for OpenBSD	19 days ago
cio		Don't provide IO when it's not set	2 months ago
cmd		ctr: fix potential panic in metric	9 days ago
containers		typo: fix misspells in comments of containers/contaiens.go	a month ago
content		Fix writer deadlock in local store	16 days ago
contrib		Bump to Go 1.11.x	6 days ago
defaults		fix words misspell	6 months ago
design		Remove outdated plugin design doc	3 months ago
diff		Set uncompressed label on diff when already exists	16 days ago



# Containerd Usage

- **Broad Usage:** containerd receives significant production usage as a core component of every release of Docker engine since early 2016. In 2018 we now see containerd used in two public cloud managed Kubernetes offerings (GKE and IBM Cloud IKS), and a long list of additional adopters found in the [TOC graduation proposal PR](#)



Firecracker



CLOUD FOUNDRY



IBM Cloud  
Kubernetes Service


PouchContainer



balena



linuxkit



Rio project

eliot



# https://github.com/cncf/toc/pull/165

## Add containerd graduation review proposal #165 Edit

Open estesp wants to merge 1 commit into `cncf:master` from `estesp:containerd-graduation`

Conversation 15 Commits 1 Checks 0 Files changed 1

Changes from all commits File filter... Jump to... +66 -0 Review changes Diff settings

66 reviews/graduation-containerd.md <> Copy path View file 🖨 ✎ 🗑 ⌵

### Containerd Graduation Proposal

Since the [March 2017 announcement](#) at KubeCon Berlin that containerd would be joining the CNCF, the containerd community has been busy continuing the development and testing of a stable, core container runtime for cloud native use cases.

In late 2017, containerd announced its [1.0 release](#), followed in early 2018 by a 1.1 release that has been updated with several minor fix releases through 2018. The containerd team has just completed the next major release, [1.2.0](#), with more complete Windows runtime support and a new runtime shim v2 API that is [proving valuable](#) for [runc](#) alternatives like [Kata containers](#).

At this time, we believe containerd is ready for the [graduation stage](#) within the CNCF, and per the guidelines listed there, we detail our readiness in the section below.

### CNCF Graduation Criteria

#### Committers from at least two organizations.

Containerd has had a variety of maintainers and reviewers since its inception, and currently have 12 committers representing Docker, NTT, Google, IBM, Microsoft, Facebook, Tesla, and Cruise Automation. We also recognize **LGTM** rights for a group we





# History of containerd

## containerd 0.2

- Integrated in Docker 1.11
- Simple runtime manager on top of runc

Early 2016

## containerd to CNCF

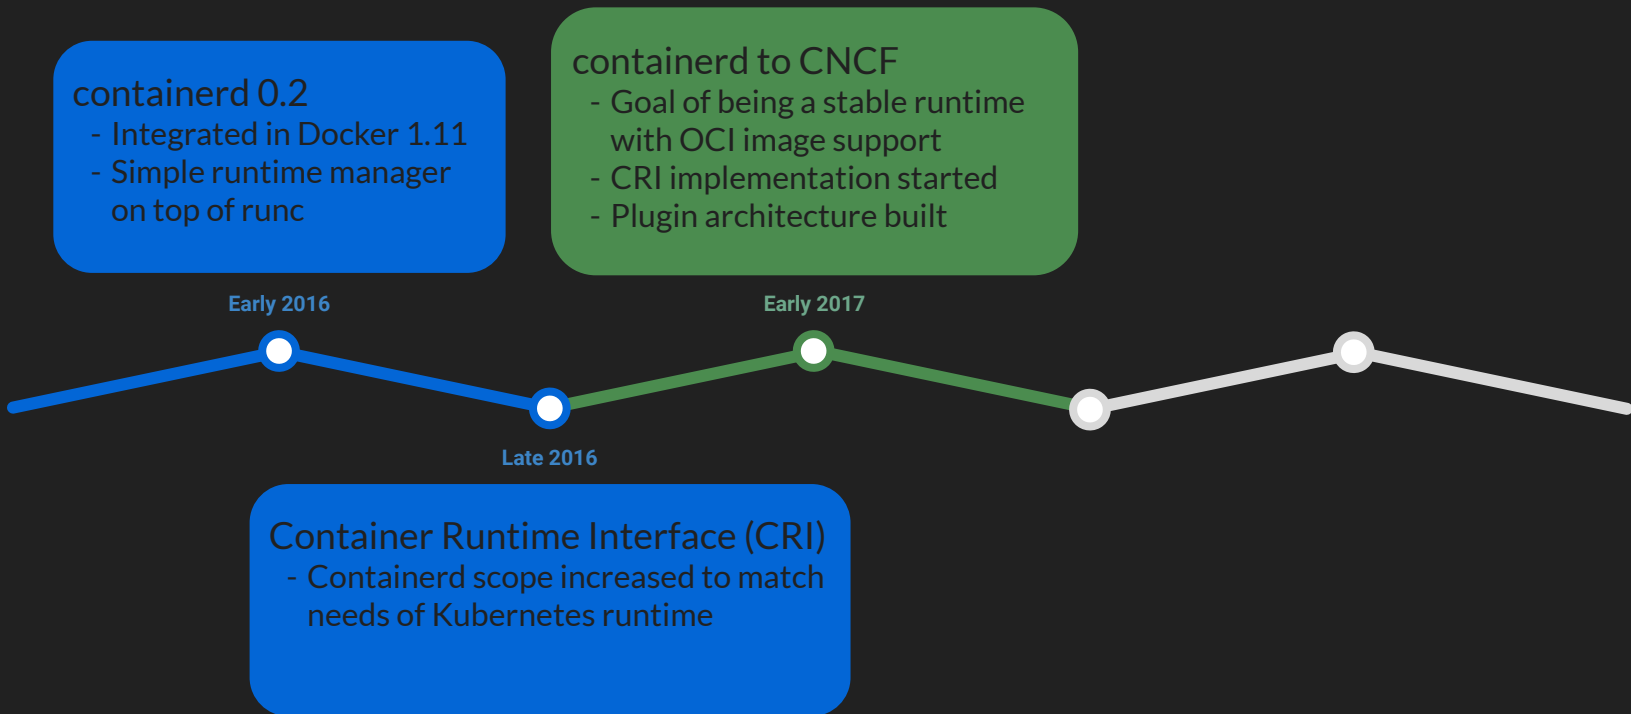
- Goal of being a stable runtime with OCI image support
- CRI implementation started
- Plugin architecture built

Early 2017

Late 2016

## Container Runtime Interface (CRI)

- Containerd scope increased to match needs of Kubernetes runtime





# History of containerd

## containerd 1.0

- Released in December '17
- API stabilized
- CRI implementation goes alpha in November

Late 2017

## containerd 1.2

- Released in October 2018
- Runtime shim API stabilized
- Focus on stability and extensibility

Late 2018

Early 2018

## Containerd 1.1

- Released in April 2018
- CRI implementation goes BETA
- CRI added to containerd as built-in plugin



Created in December 2015  
> 55 releases (v1.2.2 latest release)  
> 165 contributors  
> Docker contributed to CNCF in 2017  
> Significant and growing list of adopters

## Current Status

- Several major clouds adopting containerd as the Kubernetes CRI runtime
- Security audit performed in 4Q2018; extremely positive report of code quality and security posture
- Containerd presented to CNCF TOC for graduation in 4Q2018; all graduation criteria met; waiting on TOC vote
- Integration with and growth in broader array of use cases (HPC, AWS Firecracker VMM) proving out usefulness and capability of pluggable design and clean API





# Architecture Overview

clients






Kubelet



Container platform     
 containerd client



Prometheus

containerd

 CRI API   
CRI Plugin  
 containerd client

 containerd API   
GRPC API service

Metrics API  
Stats collector

containerd Service Interfaces

Metadata (namespaced)	Namespaces	Images	Containers	Diff	Tasks	Events
-----------------------	------------	--------	------------	------	-------	--------

Snapshotter

overlay	btrfs
native	plugin


Content Store

Runtime manager

containerd-shim

 runc	 runhcs	 kata	plugin
----------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------	--------

OS





# Smart Client Design

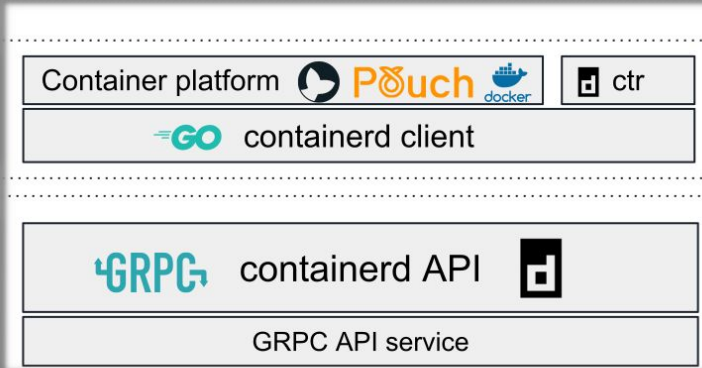


## Smart client

- Higher level interface provided by client library
- Responsible for push and pull
- Direct access to low level resources (e.g. snapshots)
- Creates container OCI configuration



# gRPC API



## gRPC API

- Low level access to components
- Mirrors internal component interfaces
- Snapshots, Content, Containers, Task, Events, etc




# Metrics

## Metric API

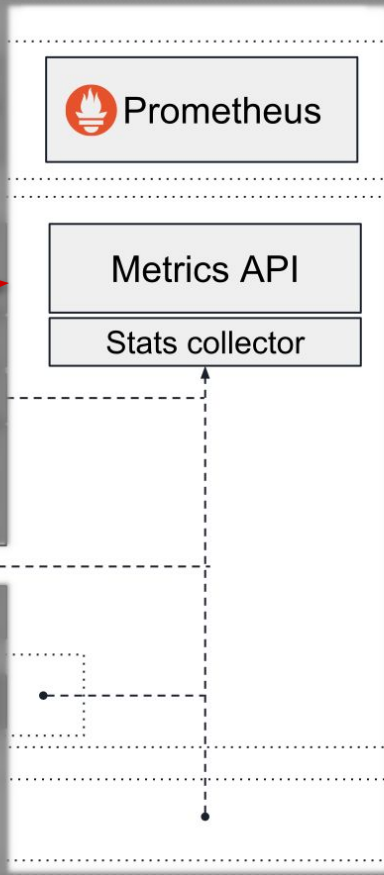
- Metrics exposed through Prometheus API
- Exposes metrics for containerd process AND container level metrics
- Enabled in containerd config `/etc/containerd/config.toml`

```
[metrics]
  address = "localhost:9090"
```



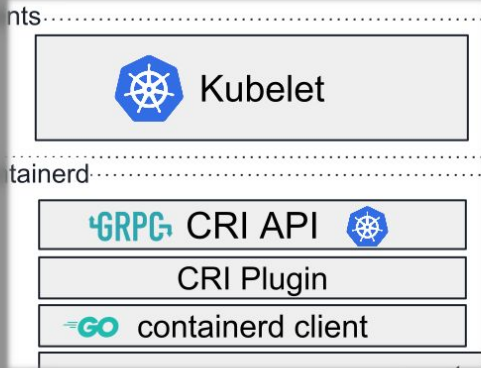
Metrics API

Stats collector





# Kubernetes Support

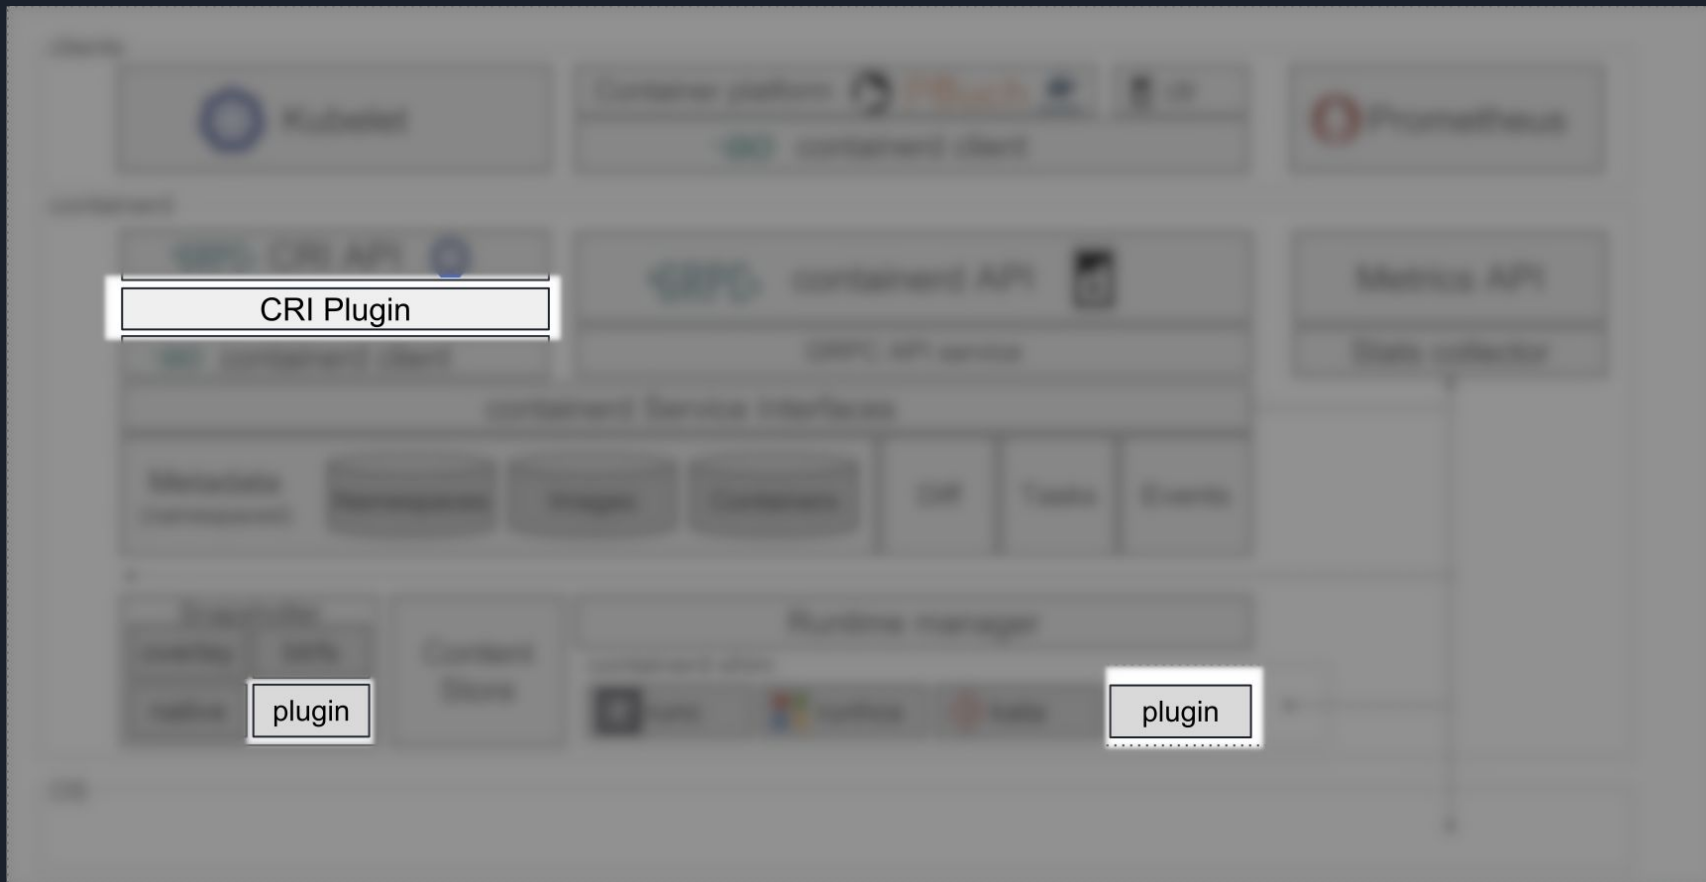


Kubernetes Runtime Support

- CRI GRPC API exposed from containerd
- Kubelet can be configured to use containerd as runtime

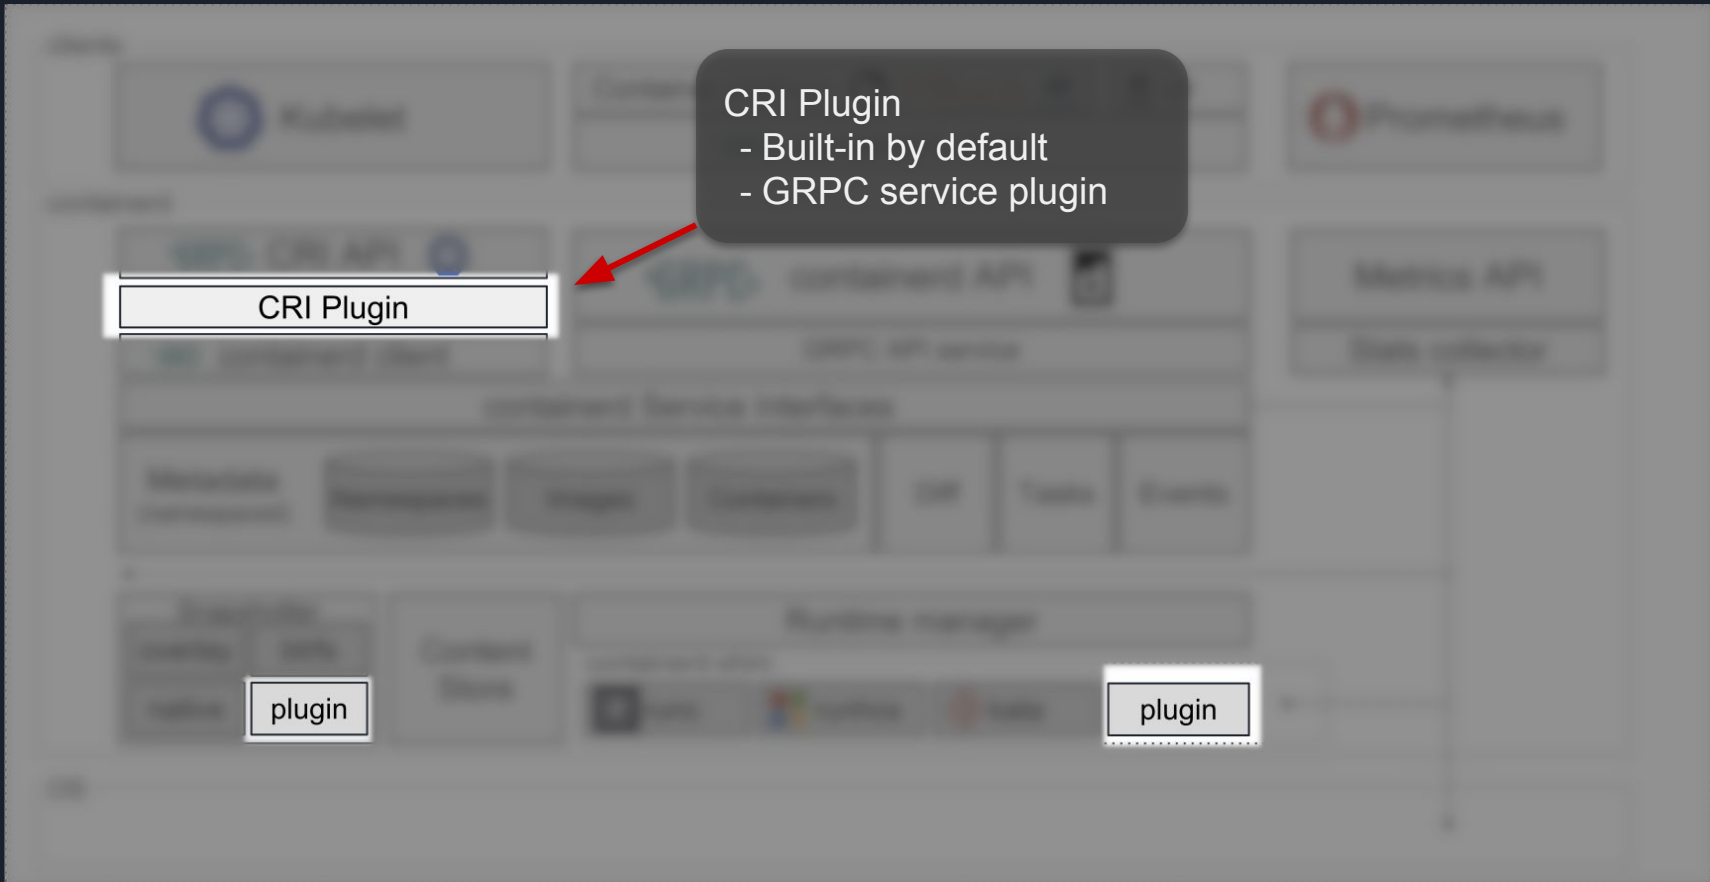


# Plugins





# Plugins (CRI)





# Plugins (Snapshotter)

CRI Plugin

Snapshotter Plugin

- Built-in (overlay, btrfs, aufs)
- Supports custom plugins over GRPC in 1.2

plugin

plugin





# Plugins (Runtime)

CRI Plugin

## Runtime plugin

- Support for custom shims in 1.2
  - Binary which implements runtime API
  - Useful for VM runtimes
- Support for OCI runtimes
- Install through `ctr install``

plugin

plugin



# More Extensibility

- **Smart client model (Golang)**
  - Resolver interface allows custom pull flow
  - Direct access to containerd interfaces
- **Server plugin architecture**
  - All services are self registered
  - Custom gRPC services
    - CRI is a gRPC plugin
  - Direct access to internal services

# containerd Adoption

- ▣ Kelsey Hightower's "[Kubernetes the Hard Way](#)" deploys **containerd** as the kubelet runtime

PouchContainer



Firecracker




IBM Cloud  
Kubernetes Service



kata  
containers



balena



CLOUD FOUNDRY

eliot



# Maturity and Integrations

## > CRI-containerd

- A gRPC plugin, calls into containerd services

## > Docker

- Use containerd's client
  - Using runtime since 17.12+
  - Using image backend in 2019

## > Moby Buildkit


- Using containerd's client
- Uses containerd interfaces directly for standalone mode

## > Alibaba Pouch

- Uses containerd 1.0.x as container runtime, integrated with image and runtime

## > Other Examples

- CloudFoundry Garden runtime
- Kata containers; AWS Firecracker
- Michael Crosby's "boss" project
- Evan Hazlett's "stellar" project



## CNCF Maturity/Status

### > Security Review

- Completed security review thanks to CNCF funding in Dec. '18
- Report is published online

### > CNCF Project Graduation

- TOC proposal created Oct '18
- <https://github.com/cncf/toc/pull/165>
- Presented to TOC Nov '18
- Expect graduation vote early 2019

# Thank you!



@estesp



github.com/estesp



estesp@gmail.com



<https://integratedcode.us>



Slack/IRC: estesp