

Ц

: COD-Working Group | Charter

: Join CNCF | Add #tag-runtime

CNCF CoC

■■: TBD (Mailing List)

Calendar invite is on:

https://www.cncf.io/community/calendar/

17: Tuesday, once in every two weeks

?: 7am Pacific (in case of US/EU winter/summer time conflict, EU time slot used)



https://zoom.us/j/94247196491?pwd=WWFCRnZFK0 RVNWZkcDVIMjZSekhZZz09

Passcode: 77777

TAG Runtime playlist

Future Meetings

Proposing Agenda Items...

- anyone can propose small-ish items
- ask/discuss with the group when to schedule
 - o may add to meetings on-the-fly (as time permits)

Pending items

•

Pending Presentations:

• Name: PR (contact name, e-mail)

Planned Meetings

Zoom:

https://zoom.us/j/94247196491?pwd=WWFCRnZFK0RVNWZkcDVIMjZSekhZZz09

Jul 22, 2025

Attendees:

Reid Priedhorsky (LANL, he/him)

Jul 8, 2025

Attendees:

• Evan Lezar (NVIDIA, he/him)

Jun 24, 2025

Attendees:

• Evan Lezar (NVIDIA, he/him)

Agenda:

No quorum, no meeting today.

Jun 10, 2025

Attendees:

Alexander Kanevskiy

No quorum, no meeting today.

May 27, 2025

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Alexander Kanevskiy
- Krisztian Litkey

Agenda:

 Docker will enable CDI by default in v28.2.0 – https://github.com/moby/moby/issues/45192#issuecomment-2898174450

May 13, 2025

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Krisztian Litkey
- Alexander Kanevskiy
- Anastassios Nanos

Agenda:

Discussion around CDI support for AMD devices:

https://github.com/moby/moby/issues/49824#issuecomment-2829013805

- Mapping --gpus flag to CDI devices makes sense.
- We discussed adding an explicit "gpu" type / signifier to the CDI spec, but ultimately felt that the device class (part of the kind in the CDI spec) was sufficient.
- Strongly suggest against using an envvar and runtime as in https://github.com/moby/moby/pull/49952 (elezar to comment / review).
- CDI support in --devices flag in nerdctl:
 - https://github.com/containerd/nerdctl/pull/2280
 - https://github.com/containerd/nerdctl/pull/4170
 - Available in v2.1.0 release: https://github.com/containerd/nerdctl/releases/tag/v2.1.0

- CDI support in Charliecloud was presented at "High Performance Software Foundation Conference" by Reid. It was well received.
- Question: Do we need a dedicated Slack channel?
 - We will continue to use the #tag-runtime slack channel at least as long as the state of tags and work groups is in flux.
- Presentation on CDI in Kata:

0

Apr 29, 2025

Attendees:

• Evan Lezar (NVIDIA, he/him)

Agenda:

 Discussion around CDI support for AMD devices: https://github.com/moby/moby/issues/49824#issuecomment-2829013805

Apr 15, 2025

Attendees:

- Alexander Kanevskiy
- Ed Bartosh
- Reid Priedhorsky (LANL, he/him)

Agenda:

- Future of the working group
 - Due to CNCF re-structuring we are looking for ways how this group will be operating in the future
 - We continue to operate in current shape until overall CNCF structure will be more understood.
- PSA: Coming features: Network Device injection
 - https://github.com/cncf-tags/container-device-interface/pull/269

0

Apr 1, 2025

Note: This is the week of kubecon, so we may not have a quorum. We will have the meeting to discuss any pressing issues, but it may be shorter than normal.

Mar 18, 2025

- Evan Lezar (NVIDIA, he/him)
- Ed Bartosh
- Alexander Kanevskiy
- Krisztian Litkey

Agenda:

- v1.0.0
 - No problems reported (YET)
 - What does this mean for the workgroup?
 - What does this mean for CDI spec and shepherdship?
 - We should reassess / update our workgroup goals. (Where are these documented?)
 - https://github.com/cncf/tag-runtime/blob/main/wgs/cod/charter/index.md
 - containerd update is required? (Krisztian Litkey)
 - v1.7.x
 - v2.0.x
 - V2.1.x
 - o cri-o (Krisztian Litkey)
 - podman / docker (Evan Lezar)

4 Mar 2025

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Ed Bartosh
- Reid Priedhorsky (LANL, he/him)
- Alexander Kanevskiy
- Krisztian Litkey

Agenda:

- v1.0.0
 - https://github.com/cncf-tags/container-device-interface/pull/257
 - Have we aligned on releasing the tip of main as v1.0.0 and specs-go/v1.0.0?
 - @elezar create a release issue template so that we don't miss any steps when releasing a new version.
 - What about the "standardized hooks" documentation PR https://github.com/cncf-tags/container-device-interface/pull/225?
 - These are based on the NVIDIA Container Toolkit and may need to be reworked for the upstream?
 - We will hold on merging this PR, and rather work on a better longer-term solution that fits better with upstream use cases and includes a reference implementation (non-NVIDIA). These can then be consumed by the NVIDIA Container Toolkit.
 - NVIDIA will move the documentation from #225 to their own documentation.
- Docker BuildKit error:
 - Why is this triggered in Docker but not on Containerd / Cri-o?

Feb 18, 2025

We don't have a meeting today: Evan is not available, no quorum.

Feb 4, 2025

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Alexander Kanevskiy
- Krisztian Litkey
- Ed Bartosh
- Reid Priedhorsky (LANL, he/him)

Agenda:

- Timeline for v1.0.0:
 - ■ COD-WG Meeting Notes
 - This should not represent additional "friction" w.r.t releasing spec updates. We can still
 use the "MinimumRequiredVersion" functionality to ensure that specs are as broadly
 applicable as possible when newer features are not used.
- Charliecloud MR:
 - https://gitlab.com/charliecloud/main/-/merge_requests/1902
 - Evan will review the changes above by the end of the week.

Jan 21, 2025

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Alexander Kanevskiy
- Ed Bartosh
- Krisztian Litkey

- Producer API:
 - Updated PR
 - elezar: update the PR description with the goals of the change
 - Split validation changes to separate PR
 - https://github.com/NVIDIA/nvidia-container-toolkit/pull/872
 - Validation:
 - Review draft PR for high-level proposal: https://github.com/cncf-tags/container-device-interface/pull/247
 - The intent is to consider pulling the validation into the producer package.
- Renaming:
 - "Device" https://github.com/cncf-tags/container-device-interface/issues/246
 - The renaming of Device is highly disruptive given how widely CDI is adopted.
 - We can improve documentation to better distinguish between different usages:
 - Top-level CDI device
 - Physical devices
 - Device nodes
 - "Spec" https://github.com/cncf-tags/container-device-interface/issues/245
 - This is modelled on the OCI Runtime Specification
 - We may be using this inconsistently and will do a pass of the documentation, error messages etc. to clarify the different usages:
 - CDI specification

- Spec
- YAML / JSON files
- OCI runtime specification
- Clarification of paths:
 - o https://github.com/cncf-tags/container-device-interface/issues/244
 - The nature of the paths are defined by their use in the OCI runtime specification.
 - We should update our documentation to indicate the allowed nature of paths.
- CDI 1.0 readiness:

https://github.com/cncf-tags/conta iner-device-interface/issues/206

Jan 7, 2025

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Ed Bartosh
- Krisztian Litkey

Agenda:

- Sorting of Mounts:
 - Should we do a point release that includes https://github.com/cncf-tags/container-device-interface/pull/241
 - o elezar will check what the history on main is and determine whether we should backport
 - The history includes changes that are targeted at v1.0.0 and should not be included in a point release.
- v1.0.0 spec:
 - Should we allow empty container edits at a device level if there are top-level edits specified?
 - Should we allow aliases?
 https://github.com/cncf-tags/container-device-interface/issues/86
 - Should we support meta devices?
 https://github.com/cncf-tags/container-device-interface/issues/71
 - Should we add support to specify required capabilities?
 https://github.com/cncf-tags/container-device-interface/issues/55

Dec 10, 2024

Attendees:

- Alexander Kanevskiy
- Reid P.
- Krisztian Litkey

Agenda:

No topics for today.

Nov 26, 2024

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Ed Bartosh
- Alexander Kanevskiy

Agenda:

- "Producer" API
 - o Is there a better name?
 - https://github.com/cncf-tags/container-device-interface/pull/233
 - Two go submodules:
 - api/validator
 - api/producer

•

Nov 12, 2024

This occurrence is canceled due to KubeCon NA 2024.

Those who can can meet in person during the conference / contributor's summit. The following members will be attending:

- Evan Lezar
- Kevin Klues
- Zvonko Kaiser
- Alexander Kanevskiy

Oct 29, 2024

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Krisztian Litkey
- Alexander Kanevskiy

- Meeting occurrence: Can we switch it to the EU timezone to avoid confusion twice a year?
 - Decision: We'll call it out in the meeting notes twice a year to reduce overhead
 - ■ COD-WG Meeting Notes
 - o Alternative, we leave it as is and accept the mismatch twice a year.
- "Producer" API
 - o Is there a better name?
 - https://github.com/cncf-tags/container-device-interface/pull/233

- CAPs/namespaces issue
 - https://github.com/cncf-tags/container-device-interface/issues/55
 - There were security concerns around modifying anything other than resources.
 - Current idea to have "require-cap" or "require-namespace" concepts and function in API to validate requirements. Will discuss further with runtimes.

Oct 15, 2024

Attendees:

- Evan Lezar (NVIDIA, he/him)
- Alexander Kanevskiy (Intel)
- Krisztian Litkey
- Ed Bartosh

Agenda:

- Rust implementation is live:
 - https://github.com/cncf-tags/container-device-interface-rs
 - Evan Lezar ping Zvonko to get updates on the status of the repo
 - dependencies updates?
 - The builds are failing?
 - Can we automate the detection of drift?
 - Add actions on both repos to encode and decode specs?
 - The most important check here would be that for a given CDI spec, the applied OCI spec modifications must be logically equivalent.
- Path to v1.0.0
 - Should we consider the maturity of the rust implementation for this milestone?

Oct 1, 2024

Attendees:

- ◆ Evan Lezar (NVIDIA, he/him)
- Alexander Kanevskiy
- Krisztian Litkey

Agenda:

- Meeting minutes document should now be publicly visible.
 - o Thanks!
- No quorum today and no urgent agenda items -> meeting is canceled for today.

Sep 17, 2024

- Evan Lezar (NVIDIA, he/him)
- Alexander Kanevskiy

- Krisztian Litkey
- Ed Bartosh
- Reid Priedhorsky (LANL, he/him)

Agenda:

- Meeting minutes are not accessible to the public:
 - o elezar will follow up to get the document public again
- Charliecloud CDI adoption still in progress
- We need to add a note on the security implications of hooks. These are typically run with root
 permissions and have access to the container and host filesystems. Steps should be taken to
 ensure that these only have access to the resources required.
- CDI in device plugin API:
 - The only thing that is required is the change of KEP status.
 - Everything was done in 1.31

Sep 3, 2024

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Krisztian Litkey
- Ed Bartosh

Agenda:

- Formal description of the update-Idcache and create-symlink hooks.
 - https://github.com/cncf-tags/container-device-interface/pull/225
 - Evan Lezar to follow up on comments on the PR
 - We may want to provide a reference implementation
 - We should update the PR to reference the implementation instead of defining the behaviour completely in the spec.
 - Decide between cobra / urfave for CLI
- Initial support in LXD has been merged:
 - https://github.com/canonical/lxd/pull/13562
 - Need to follow up to make this more generally applicable there are some NVIDIA specifics at the moment.
 - May depend on hook definitions where we could provide tooling to parse arguments, for example.
- CDI device support in NRI
 - We still want a way to provide access to CDI devices for management containers.
 - Krisztian will create a PR against 1.7 and 2.0 to bump the dependencies / wrappers.
 - Consider adding PR to v2 blockers: https://github.com/orgs/containerd/projects/9

Aug 20, 2024

- Evan Lezar
- Megan Phinney

Agenda:

- Formal description of the update-Idcache and create-symlink hooks.
 - https://github.com/cncf-tags/container-device-interface/pull/225
 - Would be good to also get feedback from Charliecloud / LXD

Aug 6, 2024

Attendees:

- Evan Lezar
- Reid Priedhorsky (LANL)
- Ed Bartosh
- Zvonko Kaiser

- Rust repository for CDI:
 - https://github.com/cncf-tags/container-device-interface/issues/150
 - https://github.com/cncf/toc/issues/1389
 - Notes:
 - Some support for the go bindings in Kata
 - Kata is moving to a complete Rust implementation
 - Also using the OCI rust implementation
 - Need a proper rust implementation for CDI
 - Current implementation is at v0.7.0
 - One option is to first develop in Kata and then donate it
 - Question:
 - Should we run in lockstep?
 - A: the primary consumer (Kata) covers different use cases than the go implementation, so lockstep may not make sense.
 - Use a separate repository instead of a sub-folder:
 - Different tooling
 - Different CI requirements
 - Currently looking at 2 maintainers:
 - Zvonko
 - Ya-an Lee
 - Function names are the same where Go and Rust idioms don't conflict.
 - It would be good to cross-check the different implementations between Rust and Go.
 - Question: Are there any other major consumers other than Kata for regular and Confidential Containers use cases.
 - Kata is the primary consumer at this stage.
 - We don't need to migrate the repo, would rather start with an single initial push.
 - Q: What is the timeline for this?
 - It would be good to get Alexander to give his OK.
 - Provenance information for the binaries?

- We don't release binaries, the tooling is consumed as go packages / modules.
- Q: Would Rust allow the Spec to be consumed separately?
 - In go, we release the spec as a separate go submodule.
 - For example: https://github.com/cncf-tags/container-device-interface/releases/tag/specs
 -qo%2Fv0.8.0
 - https://github.com/cncf-tags/container-device-interface/releases/tag/v0.8.0
- CDI support in LXD is under review:
 - https://github.com/canonical/lxd/pull/13562
- Charliecloud Implementation:
 - Is in progress including some refactoring
 - Hooks can be ambiguous around where they run
 - @elezar: will add documentation on the intent of the update-ldcache and create-symlinks hooks.

Jul 9, 2024

Attendees:

- Evan Lezar
- Mike Brown (IBM/containerd)
- Reid Priedhorsky (LANL)

- Add CDI-related presentations to README file
 - Created https://github.com/cncf-tags/container-device-interface/pull/213
- Released v0.8.0 including v0.8.0 spec
 - o https://github.com/cncf-tags/container-device-interface/releases/tag/v0.8.0
- CDI support Charliecloud (PR #1902):
 - o Something related: https://github.com/NVIDIA/nvidia-container-toolkit/issues/585
 - Regarding mount options:
 - https://github.com/NVIDIA/nvidia-container-toolkit/blob/e4cdc48854fd128b343ec2 141e61f165b946ef82/internal/discover/mounts.go#L101
 - https://github.com/NVIDIA/nvidia-container-toolkit/blob/e4cdc48854fd128b343ec2
 141e61f165b946ef82/internal/discover/ipc.go#L73
 - o CDI devices in CRI:
 - https://github.com/kubernetes/enhancements/issues/4009
 - https://github.com/kubernetes/enhancements/pull/3731
- Notes on OCI in HPC:
 - https://supercontainers.github.io/containers-wg/docs/hpc-projects/charliecloud/#detailedcritique-of-oci
 - Charliecloud papers:
 - Charliecloud: Unprivileged containers for user-defined software stacks in HPC (2017)
 - Minimizing privilege for building HPC containers (2021)
 - Charliecloud's layer-free, Git-based container build cache (2023)

Zero-consistency root emulation for unprivileged container image build (2024)

Jun 25, 2024

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Krisztian Litkey
- Ed Bartosh

Agenda:

- CDI support in LXD
 - https://github.com/canonical/lxd/pull/13562
 - elezar has made a first pass a reviewing
- CDI support in Charliecloud
 - https://github.com/hpc/charliecloud/pull/1902
- Containerd 2.0 release:
 - Currently <u>2.0.0-rc.3</u>
 - Some things that are still pending at an API level
 - Should push for nerdctl integration following this
- Presentation on CDI at WeAreDevelopers conference in Berlin:
 - o https://www.wearedevelopers.com/world-congress
- Add CDI-related presentations to README file
 - Created https://github.com/cncf-tags/container-device-interface/issues/209
 - Will ping others to add details.
- Add meeting details to README
 - Link to calendar
- Release v0.8.0
 - Required for k8s.
 - o Call out in release notes that this breaks clients that depend on the ToOCI functions.

Jun 11, 2024

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Krisztian Litkey
- Reid Priedhorsky (LANL HPC)
- Gabriel Mougard (Canonical)

- Add support for CDI devices to NRI
 - Motivation: We want to be able to trigger the injection of CDI devices (or a single device) into a container in k8s to allow for resources to be available to "management" containers.
 - Early draft PR:
 - https://github.com/containers/nri-plugins/pull/327/files

- Summary:
 - We will add native CDI support to NRI
 - We're ok to use namespaces to control access to start with
 - Evan Lezar will add more context to a README in the PR
- CDI support in LXD
 - https://github.com/canonical/lxd/pull/13562
- LANL
 - Charliecloud
 - https://github.com/hpc/charliecloud/pull/1902
 - Lightweight container implementation for HPC applications
 - Not OCI compliant
 - OCI is not really a good match for HPC applications
 - OCI doesn't map well to unprivileged environments
 - Unprivileged
 - /dev and /proc are bind-mounted
 - Devices:
 - NVIDIA GPU and Cray HPE networks
 - Working on parsing CDI spec to get a list of mounts etc.
 - Charliecloud uses user namespaces and has been using them for a long time

May 28, 2024

Attendees:

- Evan Lezar
- Alexander Kanevskiy

Agenda:

- Review of https://github.com/cncf-tags/container-device-interface/pull/207
 - Adding "stale" labels to issues and PRs.
- Review of https://github.com/cncf-tags/container-device-interface/pull/208
 - Move ToOCI functions out of specs-go
 - "Breaking change" in that these functions are no longer defined on the spec types.

May 14, 2024

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Ed Bartosh
- Krisztian Litkey

- Feature request: Extend the spec to support predefined hooks https://github.com/cncf-tags/container-device-interface/issues/203:
 - Examples:
 - Update LDCache in container

- Create Symlinks relative to container paths
- We would prefer not to do that:
 - We depend on the internals of a container. e.g. Idconfig may not be present in the container.
 - Container is supposed to be immutable.
 - Does it require the installation of a binary? The maintenance of this hook adds additional overhead.
- Windows support #205
 - Windows support is not currently the focus of the CDI specification. We are willing to work with the community to extend the specification to support Windows use cases.
 Here a real-world use case would be ideal.
- Created https://github.com/cncf-tags/container-device-interface/issues/206 to track road to v1.0.0
 - ■ CDI for HPC
 - We will continue to use CDI as the specification to define what is required in all environments (even HPC). We can provide documentation and examples on how to map from CDI specification container edits to modifications other than OCI runtime specification modifications. (e.g. mounts become bash commands).

Apr 30, 2024

Attendees:

- Evan Lezar
- Ed Bartosh
- Victor Lu

Agenda:

- Is CDI limited to accelerators?
 - No, this is not limited to accelerators.
 - What about CPU architectures?
 - The abstractions are at a higher-level
 - There is not really an intersection between CDI and low-level details such as CPU architecture.
- https://uxlfoundation.org/
- https://github.com/aep-dev
- CDI file in Kubernetes
 - o Can we import this from the CDI repo instead?
 - We may need to reorganize the packages.
 - Part of the v1.0.0 discussed two weeks ago

Apr 16, 2024

- Evan Lezar
- Ed Bartosh
- Patrick Ohly

Agenda:

- Improve OCI runtime spec dependency:
 - https://github.com/cncf-tags/container-device-interface/issues/201
- Should we release a v1.0.0 spec? What is required for this?
 - See also COD-WG Meeting Notes
 - Question: Is the spec and API stable enough?
 - The API seems stable and the spec seems to do what it needs to do for at least the DRA use cases.
 - Runtimes could still depend on older versions of the API? This places a restriction on the spec version that can be loaded by the runtimes.
 - Does it make sense to separate the spec from the API in terms of GIT repositories.
 - What is required:
 - Goal is to tag the specs-go submodule as v1.0.0.
 - Migrate the minimal version logic to specs-go package.
 - Need to check <u>golang.org/x/mod/semver</u> dependency (simplest is to check k8s and then use that version).
 - May need to switch semver packages or implement comparisons explicitly.
 - Need to document the version match process when loading specs.
 - Need to document the use of the packages / spec from the perspective of producers (DRA drivers) and consumers.
 - Better split consumption / producer.
 - Need to remove dependency on runtime-spec from spec-go should be moved to "consumption" API.
 - Need to add comments to each field in the spec to indicate when it was added.

Apr 2, 2024

Attendees:

- Evan Lezar
- Ed Bartosh
- Alexander Kanevskiy
- Christopher Desiniotis

- KubeCon EU Feedback:
- v0.7.0 Release:
 - Need to tag and update dependencies
 https://github.com/cncf-tags/container-device-interface/issues/199
 - Features:
 - Deprecation of Registry APIs
 - IntelRDT fields
 - AdditionalGIDs
- Upcoming presentation at Container Plumbing Days: https://sched.co/1aVoV
 - Do we have any example usage of CDI with other devices (Intel)

- https://github.com/intel/intel-resource-drivers-for-kubernetes
 Intel DRA driver for Intel
 GPUs
- o https://github.com/sylabs/singularity/issues/1094#issuecomment-2012264651
- o CDI is something fundamental that enables DRA use cases.
- CDI specification generation
 - NVIDIA will be looking at streamlining the generation of specifications for the interactive use case.
 - UDEV rules / some binary to generate CDI specifications
 - Does having "executable" specs make sense? Would rather be explicit rather than implicit.

Mar 19, 2024

This occurrence is canceled due to KubeCon

If you're interested, catch us at the TAG Runtime Booth!

Talks related to group:

- Sharing Is Caring: GPU Sharing and CDI in Device Plugins Evan Lezar, NVIDIA & David Porter, Google: https://sched.co/1YeQ7
- Unlocking the TAG-Runtime Magic: Where Cloud-Native, Workloads, and Al Join Forces! -Ricardo Aravena, Truera; Rajas Kakodkar, VMware; Alexander Kanevskiy, Intel; Danielle Tal, Microsoft: https://sched.co/1Yhhq
- Emerging Technologies from TAG-Runtime Working Groups | TAG Lightning Talk: https://sched.co/1aQiX

Mar 5, 2024

Attendees:

- Evan Lezar
- Krisztian Litkey
- Alexander Kanevskiy
- Antonio Ojea

- v0.7.0 Release:
 - Deprecation PR is outstanding: https://github.com/cncf-tags/container-device-interface/pull/195
 - This does not prevent / force anyone from using the existing API. If they are using tooling (e.g. golangci-lint) it may require them to mark the deprecation as ignored while they continue to use the API.
 - o elezar will create another tracking issue for the version bump.
 - https://github.com/cncf-tags/container-device-interface/pull/189 do we want to change this Configure API.
 - We will NOT merge this API change and keep Configure returning an error.
 - Containerd and CRI-O which do use this function will not need to change.

- As a follow-up (for a point release) we will address
 https://github.com/cncf-tags/container-device-interface/pull/186 to make any internal changes required.
- If we require Option to return an error at a later stage we can make this change without affecting clients (and Configure will already return an error).
- Elezar will add a comment to the PR (189) to summarize this decision

•

- Triaging of issues
 - What tooling should we enable for that?

Feb 20, 2024

Attendees:

Agenda:

- Evan is traveling today, might make sense to cancel today and chat asynchronously about item below in Slack
- Status of top-level API?
 - https://github.com/cncf-tags/container-device-interface/pull/192/files

•

Feb 6, 2024

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Ed Bartosh
- Krisztian Litkey

Agenda:

Continue chat in slack.

Jan 23, 2024

Attendees:

- Evan Lezar
- Ed Bartosh
- Alexander Kanevskiy
- Krisztian Litkey
- Markus Lehtonen

- v0.7.0 spec release:
 - Intel RDT fields:
 - Not urgent to release. Can wait pending other targets for release.
 - AdditionalGIDs:

- See concerns:
 - https://github.com/cncf-tags/container-device-interface/issues/175#issuecommen t-1885339603
- Waiting on user feedback: https://github.com/NVIDIA/nvidia-container-toolkit/issues/210
- Do we want to revert these changes for the time being and only include these (if required in v0.8.0?)
- Package-level API:
 - Krisztian will create a draft PR to discuss adding a top-level APIs for Injection,
 Refresh, and GetErrors (for example)
 - Elezar Will roll back change of NewCache signature to maintain compatibility with Docker usage.
- Update clients:
 - https://github.com/cncf-tags/container-device-interface/issues/176
 - Containerd: https://github.com/containerd/containerd/pull/9677 (now needs backport)
 - Crio: Done: https://github.com/cri-o/cri-o/pull/7526
- Docker CDI support:
 - o Docker 25 is live
 - Needs an opt-in feature flag in daemon.json: {"features": {"cdi": true}}
 - o elezar: Update documentation on CDI support in Docker
- Possible presentation at Container Plumbing Days:
 - o https://events.linuxfoundation.org/container-plumbing-days/
- KubeCon EU 2024
 - Sharing is Caring: GPU Sharing and CDI in Device Plugins

As the prevalence of AL/ML workloads running on Kubernetes increases, so too does the demand for efficient management of compute resources such as GPUs. This requires features implemented in the k8s Device Plugins that make resources consumable to end-user-applications. Additionally, the necessity of partition and resource sharing to improve utilization for cost reduction becomes critical.

In this presentation we will deep-dive into the Container Device Interface (CDI) as a new option for Device Plugin authors and the flexibility this unlocks including resource sharing for GPUs.

Starting from use cases, we will take a look under the hood at how a Device Plugin exposes GPUs and different sharing options that can be used to improve device utilization and right-sizing to workloads presented, for example time slicing, MIG, and MPS. We will discuss how k8s integrates with devices and CDI; GPU sharing mechanisms, and how applications and frameworks can integrate with this functionality.

Jan 9, 2024

- Alexander Kanevskiy
- Ed Bartosh

- Markus Lehtonen
- Evan Lezar

Agenda:

- PR#164 OCI RDT field to CDI spec
 - Evan will point out the minimum version logic
 - Let's be consistent with envvars and mounts (it seems to be new overwrites existing)
- Additional GIDs work will be done in the next week to also be included in the 0.7.0 spec bump.
- CDI for the Device Plugin API
 - o Kevin to update the issue https://github.com/kubernetes/enhancements/issues/4009
 - O What else is required?
 - Two vendors:
 - NVIDIA GPU Device Plugin
 - Intel?
 - Reach out to Swati or Francesco for guidance on what else is required.

Dec 26, 2023

Canceled: holiday season. Let's continue the series in January 2024.

Dec 12, 2023

Attendees:

- Alexander Kanevskiy
- · Abhishek India, student
- Linda Zhou

Agenda

- How to contribute to this group
 - https://github.com/cncf-tags/container-device-interface/ -> issues, PRs, ideas, etc.

•

Nov 28, 2023

Canceled: Evan and Alexander are not available today, no urgent or proposed agenda items.

Nov 14, 2023

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Kevin Klues
- Riaan Kleinhans
- Krisztian Litkey
- Linda Zhou

- New attendeesP:
 - Riaan, from CNCF as TAG & WG PM support resource
- v0.6.2 release: https://github.com/cncf-tags/container-device-interface/issues/176

- Includes usage of tags.cncf.io
 - Podman update
 - Docker updated
 - Singularity updated
 - Cri-o is in progress, diamond dependencies due to <u>https://github.com/containers/common</u>
 - Al elezar: create issue for new release in containers/common including
- Feature proposal: https://github.com/cncf-tags/container-device-interface/issues/175
 - Should we add a check for something like gid 0?
 - Elezar will put together a PR
- Cncf-tags membership:
 - Ed is not a member of the organization
 - o I added him to the repo as well?
- KubeCon overview/feedback?
 - Discussed at the runtime meeting
 - Edge working group: (Andy and Joel were present):
 - **■** Edge-Native Application Design Behaviors
 - Special purpose OSs: Sean McGinnis presented
 - We probably need to reachout to Edge WG to talk about DRA
 - Let's join one of the WG sessions
- Will use k8s wg-batch to further discuss DRA
 - See
 - https://groups.google.com/a/kubernetes.io/g/dev/c/BDtCFfXQbw0/m/1smc98ArCAAJ
 - See agenda Kubernetes Batch Working Group Agenda

Oct 31, 2023

Attendees:

- Evan Lezar
- Krisztian Litkey
- Alexander Kanevskiy
- Linda Zhou

- Tagging v0.6.2
 - It does not seem as if a `specs-go/v0.6.0` tag is required
 - Does require that clients explicitly specify a specs-go version (e.g. v0.6.0)
 - Does this make sense:
 - https://github.com/cncf-tags/container-device-interface/pull/173/files
 - Main clients:
 - Containerd ed / evan will check containerd update
 - Crio krisztian will check the crio update
 - Podman elezar will look into this
 - Docker elezar will look into this
- Vanity URL
 - https://github.com/cncf-tags/container-device-interface/pull/172

- Renames modules to tags.cncf.io/
- o Will require that clients be updated:
 - Containerd
 - Crio
 - Podman
 - Docker
- Question: Do we want to do this first and then tag v0.6.2
- We will:
 - o Review and merge https://github.com/cncf-tags/container-device-interface/pull/173
 - Rebase, review and merge <u>https://github.com/cncf-tags/container-device-interface/pull/172</u>
 - o Tag:
 - v0.6.2
 - specs-go/v0.6.0
 - Create standalone client that can be used to test releases / imports
 - o Check on clients who have dependabot enabled (see above for list):
 - Podman (elezar)

Notes:

We will maintain backward compatibility w.r.t spec versions. A particular package version will
always read older specs. We provide a utility to ensure that the correct minimum spec version is
specified for a spec depending on the features used.

Oct 17, 2023

Attendees:

- Evan Lezar
- Ed Bartosh
- Krisztian Litkey

- CDI for Device Plugin Beta:
 - Ed Bartosh will add e2e tests
 - Date for code changes is 31 October / 1 November
- Vanity URL is in progress
- Changing injection semantics to allow for failed resolution:
 - https://github.com/cncf-tags/container-device-interface/issues/162
- Detection support for CDI in runtimes
 - This does not seem to fit in kubeadm
 - Mav extend the kubelet
 - Kubelet will guery the runtime to see if CDI is supported
 - Should log an error or trigger an error event
 - Is it possible to fail the kublet on an invalid configuration?
 - Note that there are out-of-band mechanisms to enable CDI in the runtime e.g. the NVIDIA Container Runtime
 - It may be an option to enable stricter checks.
 - Ed will create an issue

Oct 3, 2023

Agenda

- If it will be no other proposed agenda items before the meeting, consider it canceled
 - Due to public holiday in Germany, low participation is expected

Sep 19, 2023

Attendees

- Evan Lezar
- Ed Bartosh
- Alexander Kanevskiy
- Thomas Milox
- Kevin Klues

Agenda

- Intro: Thomas: active DRA contributor
- Repo migration follow-ups?
 - Vanity URL?
 - https://github.com/cncf-tags/container-device-interface/issues/156
 - https://github.com/cncf/toc/issues/1152
 - Rename modules (wait for vanity URL)
 - Documentation:
 - We can make some of the changes now, but it might be simpler to wait for the urls and module rename.
 - Has the repo been onboarded to the automatic permission tooling etc. (CLOwarden)
- Containerd 1.6.x backport follow-up
 - We will not push for backporting to 1.6.x unless there is a specific customer that requires this.
 - The Containerd maintainers are open to accepting a backport **IF** we can demonstrate an explicit customer requirement. This includes, for example:
 - DRA is GA in Kubernetes where the feature and API group are enabled by default and users are still "stuck" on 1.6.x
 - Users want CDI support in the device plugin.
 - LTS in the case of Containerd 1.6 seems to mean that it will be supported until the Containerd 2.0 release.
 - Big cloud providers have already migrated to 1.7.x
- CDI support in device plugin API:
 - Does it make sense to promote this to Beta? Yes, but we still need to spend cycles on adding end2end tests.
- DRA was not included in sig-node planning:
 - May not have been clear that we're aiming to graduate to Beta in the 1.29 release.

Action Items:

- Evan Lezar:
 - Follow up on repo migration (documentation and permission tooling)

Aug 22, 2023

Attendees

- Evan Lezar
- Alexander Kanevskiy
- Ed Bartosh

Agenda:

- Path to "stable" specification and API (triggered by <u>this Mike Brown's comment</u> in the backport CDI to Containerd 1.6 PR)
 - o (also <u>discussed</u> on May 16, 2023)
 - What is the path to 1.0.0?
 - Is CDI stable? We can call what we have now beta.
 - Elezar: I would consider the spec quite stable.
 - Getting DRA into Beta is a prerequisite for 1.0.0 since this would provide a lot of feedback on real-world use cases.
- Moving repositories to cncf-tags org (follow-up):
 - https://github.com/cncf/toc/issues/1142
 - o Repo has been transferred.
 - Al: elezar to follow up regarding prow automation / integration.
 - Still waiting on permissions to be able to merge

Aug 8, 2023

Attendees

- Evan Lezar
- Alexander Kanevskiy
- Krisztian Litkey
- Kevin Klues

- Moving repositories to cncf-tags org
 - https://github.com/container-orchestrated-devices/container-device-interface
 https://github.com/cncf-tags/container-device-interface
 - Option 1. Create new repo at cncf-tag and archive container-device-interface
 - Option 2. GitHub could move the repo:
 - Does golang also redirect? (the import paths should still work due to the redirect)
 - Follow ups:
 - Update go module name? (This may break clients that don't depend on specific versions). The update is simple enough using a replace rule
 - Should we separate the spec and the utils as part of this? (If we just migrate)

- Let's first just migrate
- Create rust-specific repo
- Kata devs can start the work in their own repo and we can move / migrate that to cncf-tags once that is ready
- Evan Lezar DE: follow steps here
 https://github.com/cncf/toc/blob/main/tags/cncf-tags-github-org.md#creating-a-new-repo
 - Created:
 - https://github.com/container-orchestrated-devices/container-device-interface/pull/151 to add the missing files as described by the process.
 - https://github.com/cncf/toc/issues/1142 to request a new repository.
- Rust version of utils: #150
 - OCI Runtime spec:
 - https://github.com/opencontainers/runtime-spec/tree/main
 - https://github.com/containers/oci-spec-rs
 - What about:
 - cncf-tag/cdi-spec/
 - specs-go
 - specs-rs
 - cncf-tag/container-device-interface/pkg/cdi (imports cdi-spec/specs-go)
- Making resources visible to pod sandbox at creation:
 - Need to do all prepare resource calls before creation of pod.
 - Currently this may already be the case since resources need to be set up before CNI is enabled.
 - Needs to be added to the following KEP: https://github.com/kubernetes/enhancements/pull/4113

July: vacations. Next meeting on August 8th

Jun 27, 2023

Attendees

- Alexander Kanevskiy
- Evan Lezar

- Summer vacations:
 - Planning to cancel the summer sessions, meaning that the next meeting will be on Aug 8, 2023.
 - o If there are urgent issues, use the #tag-runtime channel in the CNCF slack
 - Evan will add a note to TAG Runtime chairs
- cncf-tags discussion:
 - https://github.com/cncf/toc/issues/1098

Jun 13, 2023

Attendees

Alexander Kanevskiy

Agenda

• No agenda and no quorum -> meeting canceled for this week.

May 30, 2023

Attendees

- Evan Lezar
- Mike Brown
- Alexander Kanevskiy
- Eduard Bartosh
- Krisztian Litkey
- Jukka Rissanen

Agenda

- Status of other CLI clients
 - o ctr created #142
 - PR <u>https://github.com/containerd/containerd/pull/8525</u>
 - Elezar: Add a summary of our discussion around local vs remote and device nodes.
 - o nerdctl create #143
 - Docker:
 - Moby PR merged: https://github.com/moby/moby/pull/45134
 - CLI PR in review: https://github.com/docker/cli/pull/4084

The Moby changes should already support the following compose spec (reference):

```
services:
 test:
    image: application-image
    command: do-the-thing
    deploy:
      resources:
        reservations:
          devices:
            # NVIDIA Devices
            - driver: cdi
              deviceIDs:
                - nvidia.com/gpu=0
                nvidia.com/gds=all
                nvidia.com/mofed=all
            # Intel Devices
            - driver: cdi
              deviceIDs:
                - intel.com/vf=0
```

- Adding CDI devices to Device Plugin API
 - https://github.com/kubernetes/enhancements/issues/4009
- Spec bump:
 - Should we release 0.6.0?
 - Annotation support
 - Allow . in class names
- Backport CDI to Containerd 1.6
 - Once it's "stable" and tested in 1.7 it should be ok to backport to 1.6.
- Behavior of Exec in modified containers:
 - At least in containerd it is elezar's experience that environment variable modifications are not visible when exec'ing to a container.
 - To be verified for cri-o and containerd

May 16, 2023

Attendees

- Evan Lezar
- Jukka Rissanen
- Krisztian Litkev
- Ed Bartosh
- Alexander Kanevskiy

- Container Device Interface (CDI) in HPC:
 - Singularity CE Community Meeting May 4th, 2023
 - HPC Containers Advisory Meeting Notes
 - Started CDI for HPC to capture thoughts
- Question was raised in HPC Containers Advisory Council and Singularity discussions
 - Have we given thought to the requirements for moving the specification to v1.0.0?
 - Proposed requirements / exit criteria for the COD working group:
 - Specification 1.0.0
 - Adoption in Docker, Containerd, Cri-o, Singularity, nerdctl, ctr
 - Hosting of GitHub organization (CNCF / k8s sigs)
 - Dynamic Resource Allocation is in GA (will require a v1.x CDI specification)
 - TODO: Revisit this in a future meeting
- CDI Devices in the device plugin API:
 - Issue: https://github.com/kubernetes/enhancements/issues/4009
 - KEP: https://github.com/kubernetes/enhancements/pull/4011
 - O WIP Branch:
 - https://github.com/kubernetes/kubernetes/compare/master...elezar:kubernetes:4009/add-cdi-devices-to-device-plugin?expand=1
- CDIDevices field in CRI support to Containerd and Cri-o
 - Containerd
 - Done

- To backported to 1.7
- Ideally to 1.6.x
- o Cri-o
- Common code collected in container-orchestrated-devices
- RDMA fields for container spec
 - Should these be added to the CDI specification?
 - o https://github.com/opencontainers/runtime-spec/blob/main/config-linux.md#rdma
 - TODO: Evan Lezar reach out to Moshe / Zvonko
- Status of other CLI clients
 - o ctr. nerdctl
 - Docker still in progress
 - Create issues for these

May 2, 2023

Attendees

- Evan Lezar
- Alexander Kanevskiy

Agenda

- PR to update permissions for spec files to 644 over 600 which is the default for CreateTemp
 - Should this be an option instead?
 - o May be better to have a producers control this instead
- Will ping Adrian for examples on other PR
- Evan Lezar will present at Singularity Community Call and HPC Advisory Council on May 4, 2023
- Docker PRs are still in review and Evan Lezar will continue interacting with the community

Mar 21, 2023

Attendees

- Evan Lezar
- Jukka Rissanen
- Kevin Klues
- Alexander Kanevskiy
- Patrick Ohly

- CDI Support in Moby / Docker
 - PRs ready for review:
 - https://github.com/moby/moby/pull/45134
 - https://github.com/docker/cli/pull/4084
 - o Includes vendor-agnostic example in PR description
- CDI Support in singularity
 - PR in draft: https://github.com/sylabs/singularity/pull/1459
 - Community call on May 4, 2023 that will demo the feature. Elezar will present short CDI overview.

- Consumer updates:
 - All updated to v0.5.4
- Annotations:
 - Annotations PR has been merged
 - Are there any other changes that we would like to get in before bumping the spec version to 0.6.0?
 - Klueska: It would be good to be able to handle conflicts better when devices have the same name
 - Could the following deduplication logic be useful: https://gitlab.com/nvidia/container-toolkit/container-toolkit/-/merge_requests/345
- Dependabot config for repo
 - Will clean up change and limit updates to runtime spec since here our concern is missing fields, for example.
 - As a library we should not force dependency updates on consumers
- Should we add CRI.CDIDevice equivalent to the device plugin's allocate call?
 - o Should do this for 1.28
 - Will require a KEP for the extension
 - Could this be a patch to an existing device plugin KEP?

Notes:

•

Mar 7, 2023

Attendees

- Evan Lezar
- Jukka Rissanen
- Kevin Klues
- Krisztian Litkey

- CDI Support in Singularity (Sylabs)
 - Meeting last week went well
 - o CDI will be used in OCI mode for the upcoming 4.0 release
 - https://github.com/sylabs/singularity/issues/1394
 - Discussion is starting on usage in native mode:
 - https://github.com/sylabs/singularity/issues/1395
 - Evan Lezar will attend a community call on Apr 6, 2023 (tentative) to discuss further
- Consumer updates:
 - Containerd has been <u>updated</u>
 - Podman was already updated
 - Crio is still at 0.5.3
 - https://github.com/cri-o/cri-o/pull/6x`655 failed and they want to update manually.
- CDI Support in Moby / Docker
 - Work no started, still on target to start by mid to end March
- Runtime tools dependencies:

- https://github.com/opencontainers/runtime-tools/issues/758
- New person who will look at some of these issues.
- Would it be feasible to split out the generator?
- Notes:
 - As we start to focus on non-OCI (and non-containerized applications) such as HPC schedulers, it may make sense to be more careful about naming.

Feb 21, 2023

Attendees:

- Evan Lezar
- Victor Lu
- Ed Bartosh
- Jukka Rissanen
- Alexander Kanevskiy
- Krisztian Litkey
- Kevin Klues
- Patrick Ohly

Agenda:

- CDI support in Singularity
 - Meeting with developers next week Feb 28, 2023 at 15:00 Berlin
 - https://github.com/sylabs/singularity/issues/813
 - https://github.com/sylabs/singularity/issues/1094
 - o I will ask whether this can be recorded and shared with other CDI maintainers
 - Questions:
 - Is this daemonless (e.g. podman)?
- Elezar still to check whether the consumers have updated CDI to 0.5.4
 - Ed will create PR for containerd
 - Krisztian will create PR for cri-o
 - Elezar will create PR for podman Updated by dependabot (https://github.com/containers/podman/blob/main/go.mod#L11)
- CDI support in Moby / Docker CLI
 - Work targeted to start at mid to end march
 - Will follow-up with moby team on release schedule (when will the feature be available?)

Feb 7, 2023

- Evan Lezar
- Ed Bartosh
- Jukka Rissanen
- Ian Forbes
- Kevin Klues
- Alexander Kanevskiy

Agenda:

- CDI support in Docker / Moby
 - Moby team is onboard to review code and integrate CDI into Moby and Docker CLI
 - Point of contact is
- NVIDIA CDI Tooling
 - NVIDIA Container Toolkit v1.12.0 is GA and includes tooling to generate CDI specs
 - o nvidia-ctk cdi generate
 - Expand examples in README
 - Working on adding support for CDI spec generation for WSL2 systems
- Cleanup "old" API in repo
 - o OLD:

https://github.com/container-orchestrated-devices/container-device-interface/tree/main/pkg

NEW:

https://github.com/container-orchestrated-devices/container-device-interface/tree/main/pkg/cdi

- o elezar will remove the files from '/pkg' and add note to the readme.
- Do a quick search of main consumers to see if anyone is using the "old" APIs (cri-o, contianerd, podman)
- Update containerd to newest CDI tagged version
 - v0.5.4 has been tagged: https://github.com/container-orchestrated-devices/container-device-interface/releases/ta
 https://github.com/container-orchestrated-devices/container-device-interface/releases/tag/v0.5.4
 - Ed will update contianerd
 - Elezar will check:
 - Podman
 - created and failing
 - https://github.com/containers/podman/pull/17375
 - Cri-o
 - not created yet
- Heads-up for everyone on the call: we got contacted by TAG-Runtime chair to update on WG activities, reports for TOC.
- Question Regarding NRI
 - o Is it possible to get the pod namespace from an NRI plugin
 - Yes, the pod information is available.
- Update Golang version:
 - o 1.20 adds something that could be used as a replacement for muliterror
 - Would be good for runtime-tools, but we may not be able to bump runtime-tools as aggressively
 - Nothing is blocking runtime-tools from being updated (currently at 1.16)
 - o Krisztian will look at creating an update to runtime tools to a newer version
 - Podman is at 1.18
 - Containerd is at 1.18 but is using 1.19 in tooling
 - Crio is at 1.19
 - Elezar will update NVIDIA deps to a newer version
- Annotations in CDI specification

Elezar will update annotations PR
 https://github.com/container-orchestrated-devices/container-device-interface/pull/85 with minimum version checks and improved / simplified validation by next meeting.

Jan 24, 2023

Attendees:

- Evan Lezar
- Ed Bartosh
- Jukka Rissanen
- Alexander Kanevskiy
- Kevin Klues
- Krisztian Litkey
- Jukka Rissanen
- Zvonko Kaiser
- Victor Lu

- CDI support in Docker:
 - https://github.com/docker/cli/issues/3864
 - Community meeting on Jan 26, 2023
 - o Evan, Kevin, and others will attend.
 - Kevin will share meeting details.
- Annotations in CDI
 - Evan will review
 https://github.com/container-orchestrated-devices/container-device-interface/pull/85
 again to get it merged:
 - Add "minimum version for annotations"
 - Need to ensure that consumers (reader) are updated to pull in new dependency
 - o PR to Kata
 - Containerd shim in Kata prepares the OCI spec to pass to Kata
 - Meta information is need to create, for example PCI Express topology, in the sandbox
 - o Problems:
 - CDI devices are only present in container annotations. This will be removed once we move to a field in the CRI create container request. (as opposed to the "run pod sandbox" request).
 - In the Kata use case, this information is available to the containerd-shim
 - Does this only work if the first container requests devices? This works for single containers. For multiple containers, this has the same issue as for memory and CPU.
 - Questions:
 - The only time at which the CDI spec is processed is to create the config file for the container. At this point the sandbox is already created.
 - Does the proposal mean that we're trying to process the spec twice? When are we processing the CDI information
 - A: This is similar to the "prestart" hook. Sometimes we need this in an outer runtime vs inner runtime.

- Do we not need a CRI message for "prepare sandbox" which can include this CDI information. This needs a KEP on Kubernetes side.
 - There is a draft of a KEP but it is not ready for submission
- Update for CRI structures:
 - o https://github.com/kubernetes/enhancements/pull/3731
- Utility for getting minimum spec version has been merged and is being consumed from NVIDIA tooling
 - What is required for https://github.com/kubernetes/kubernetes/issues/113831?
- Housekeeping:
 - Is https://github.com/container-orchestrated-devices/container-device-interface/issues/90 resolved? (PR has been merged)
 - https://github.com/container-orchestrated-devices/container-device-interface/pull/103
 - This reduces our dependencies as this doesn't pull on cobra and force this on our consumers
 - Does using `flag` make more sense?
 - This is out of scope for the time being.
 - o https://github.com/container-orchestrated-devices/container-device-interface/issues/104
 - We should respond to this AI: Alex
 - Can we use DRA for CNI interactions instead?

Jan 10, 2023

Attendees:

- Evan Lezar
- Ian Forbes
- Kevin Klues
- Krisztian Litkey
- Alexander Kanevskiy
- Ed Bartosh

Agenda:

- Review of PR for determining minimum spec version:
 - https://github.com/container-orchestrated-devices/container-device-interface/pull/93
 - Also relevant to https://github.com/kubernetes/kubernetes/issues/113831
- Annotations will ping Zvonko again to get him to commit to a meeting occurrence
- Beta acceptance criteria for DRA
 - Most of the issues are tracked on the following board: https://github.com/orgs/kubernetes/projects/95/views/1

Dec 13, 2022

- Discussion on CDI annotations
- No other topics, meeting is canceled.

Nov 29, 2022

Attendees:

- Evan Lezar
- Ian Forbes
- Alexander Kanevskiy
- Krisztian Litkey

Agenda:

- Feedback from crio meeting:
 - https://github.com/opencontainers/runtime-tools/issues/758
 - Waiting for feedback from maintainers
 - gojsonschema will be the most problematic
 - Internally we still plan to split our repos / packages by use case (producer vs consumer vs validation)
- Permission issue with device nodes not rooted at /dev in crun:
 - User created issue: https://github.com/containers/crun/issues/1047
 - Fix PR: https://github.com/containers/crun/pull/1051
- Review of PR for determining minimum spec version:
 - o https://github.com/container-orchestrated-devices/container-device-interface/pull/93
- Review of PR for adding annotations to spec and devices:
 - There are some questions about validation (since we don't actually trigger spec validation when loading / creating the spec).
 - Would need review by a non-NVIDIA contributor
 - elezar: Will ask Zvonko to present the end-to-end flow for annotation usage and how this integrates with runtimes (containerd, cri-o) and e.g. kata.
 - How will the current integration code have to be modified to enable this use case?
 - See PRs that inject devices into pods, for example.
- Question: Is everything ready for DRA in alpha state?
 - Can I check out the tip of k8s and containerd and start writing my plugins (resource drivers) using DRA
 - Code for NVIDIA resource driver has been updated to work with the code at tip
 - Will create an example resource driver repo for general use (will be based on GPU resource driver). See https://github.com/kubernetes/org/issues/3837
 - Will be ready as an alpha feature if feature-gate is enabled
 - Containerd 1.7.0 will include support for CDI (<u>1.7.0-beta.0</u> has been released)
 - Documentation and blog post is almost ready

Nov 15, 2022

- Patrick Ohly
- Ed Bartosh
- Krisztian Litkey

- Alexander Kanevskiy
- Jukka Rissanen
- Evan Lezar
- Kevin Klues

- Patrick: dependency and version control
 - As part of final reviews for DRA, @dims and @liggit pointed out that some code depends on old deps. Even where they are not used, but show up in the dependency trees. Mostly about json validation.
 - Issue:
 - https://github.com/kubernetes/kubernetes/pull/111023#issuecomment-130927372
 8
 - https://github.com/kubernetes/kubernetes/pull/111023#issuecomment-130905868
 0
 - Proposal
 - We will split our functionality into submodules / different repos as required to reduce the dependencies for pulling in only the spec, for example or using the CLI.
 - We will evaluate our exposed APIs (in the pkg/cdi package) whether these need to split according to functionality. This will be based on the DRA use cases that we have now gathered and the actual usage by clients such as containerd, crio, and podman
 - We will approach runtime-tools maintainers with suggestions / requests for feedback on how to address the concerns raised in the k8s PR.
 - Who are the maintainers? (https://github.com/kolyshkin as points of contact)
 - Raise this at the next cri-o meeting (Krisztian)
- Kevin: lifehack for github and long PRs:

```
let tryAttempts = 0;
function loadComments () {
         let needRescheduling = false;
         const buttons = document.querySelectorAll(".ajax-pagination-btn[data-disable-with]")
         buttons.forEach((button) => {
                  button.click();
                  needRescheduling = true;
                  tryAttempts = 0;
         })
         if (needRescheduling || tryAttempts < 5) {</pre>
                  if (needRescheduling) {
                            console.log("Loading comments.")
                  } else {
                            console.log("Looking for more to load.");
                  tryAttempts++;
                  setTimeout(loadComments, 500)
         } else {
                  console.log("All comments loaded.");
                  const resolvedButtons = document.querySelectorAll(".js-toggle-outdated-comments[data-view-component]");
                  resolvedButtons.forEach((button) => {
                            button.click();
                  })
                  console.log("All resolved comments loaded.")
         }
```

Nov 1, 2022

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Kevin Klues
- Ian Forbes
- Jukka Rissanen
- Chris Desiniotis

Agenda:

- Patrick: issue #88, error wrapping via native go implementation
 - Makes sense to do, will be done by Intel people
- New CDI release
 - After errorwrap fixed (#88)
 - Also requires #87
 - If Jukka / Ed cannot do this, then Evan will have a look
 - Ideally before code freeze for k8s (next Tuesday)

•

Oct 18, 2022

Attendees:

- Kevin Klues
- Evan Lezar
- Alexander Kanevskiy
- Krisztian Litkey
- Jukka Rissanen

Agenda:

- Opens
 - Podman and rootless / user namespace issue with CDI devices.
 - device: /dev/nvidia-caps/nvidia-caps255
 - https://gitlab.com/nvidia/container-toolkit/container-toolkit/-/issues/8
 - Works in runc, not in crun
 - Podman addDevice:

https://github.com/containers/podman/blob/25502d53bcda27abe753b6a2152f4ddacae42924/pkg/specgen/generate/config_linux.go#L126

• If "rootless" is detected, the path is bind mounted. This does not set up the intermediate folder (/dev/nvidia-caps)

- runc creates parent:
 - https://github.com/opencontainers/runc/blob/70e3b757c0eb39221aa5a4d beb723c88f348f5cd/libcontainer/rootfs linux.go#L683
- To check:
 - See whether cgroups are set up (major/minor)

Oct 4, 2022

Attendees:

- Ed Bartosh
- Alexander Kanevskiy
- Krisztian Litkey

Agenda:

No opens or planned agenda items: canceling.

Sep 20, 2022

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Ian Forbes
- Ed Bartosh
- Jukka Rissanen
- Krisztian Litkey

Agenda:

- Opens
 - o New release?
 - Ed: Non-Linux platform, compilation issues
 - Krisztian: to fix population of missing device on non-linux
 - New release + update to runtimes afterwards
 - o Ed: configuration?
 - Krisztian:
 - o lan: cdi + ctr
 - Ed: there is old PR, it was rejected by containerd maintainers
 - ctr PR: https://github.com/containerd/containerd/pull/6204
 - ctr PR #2: https://github.com/containerd/containerd/pull/7399

Sep 5, 2022

- Alexander Kanevskiy
- Ed Bartosh

- Krisztian Litkey
- Kevin Klues
- Jukka Rissanen
- Rishit Dagli

- Data race condition reported after containerd changes:
 - https://github.com/container-orchestrated-devices/container-device-interface/issues/81
 - Merged
 - o New tag is needed probably to update references in containerd/cri-o.
- CDI & Kata
 - https://github.com/containerd/containerd/pull/7331/
 - Krisztian will work with Fabiano (Kata) to clarify and verify case, and see if Kata require some special treatment.
- Ed: next containerd release?
 - Usually by the end of the year. No branch for 1.7 yet.

Aug 23, 2022

Attendees:

- Krisztian Litkey
- Patrick Ohly
- Kevin Klues
- Alexander Kanevskiy

Meeting is canceled today due to empty agenda and some people are not available.

Aug 9, 2022

Attendees:

- Evan Lezar
- Kevin Klues
- Krisztian Litkey
- Jukka Rissanen
- Alexander Kanevskiy

Agenda

- PR #73
- Implementation for API that creates temp files that are globally unique when creating spec files
 - Two functions
- CNCF Calendar Event (we will send a new invite): AR Sasha to ping CNCF staff again.

Notes:

- Question regarding generation of annotations: Is it possible to generate a unique key for the
 annotations instead of requiring that the user generates the name instead of allowing the API to
 do so.
 - Answer: This was related to the old device plugin where this was called multiple times.
 This may have been required for the early prototyping and not for the new DRA implementation.
 - We will rename the `deviceID` parameter for annotation to something like `suffix` to reduce confustion. The length limit will be called out in the docstring.

Jul 26, 2022

Attendees:

- Evan Lezar
- Kevin Klues

- Questions regarding CDI and Dynamic Resource Allocation:
 - Number be the first character of the device name
 - Fixed upstream; not in containerd / cri-o
 - Limit to length of fully-qualified device name (around 66 characters)
 - A: I think this may be limitation of the annotations used to select the device
 - See <u>here</u>

```
cdiDeviceName := "device-" + string(req.ClaimUid)[0:8]
cdiFilePath := filepath.Join(cdiRoot, cdiDeviceName+".json")
containerEdits := &cdiapi.ContainerEdits{}
containerEdits.Append(&cdiapi.ContainerEdits{ContainerEdits:
&specs.ContainerEdits{
        Env: []string{"NVIDIA VISIBLE DEVICES=0,1"},
    }})
    contents := &specs.Spec{
        Version: cdiVersion,
        Kind: cdiKind,
        Devices: []specs.Device{{
            Name:
                           cdiDeviceName,
            ContainerEdits: *containerEdits.ContainerEdits,
        }},
    }
spec, err := cdiapi.NewSpec(contents, cdiFilePath, 0)
//return qualified name of the cdi device
dev := spec.GetDevice(cdiDeviceName).GetQualifiedName()
```

```
klog.Infof("Return: " + dev)
return &drapbv1.NodePrepareResourceResponse{CdiDevice: []string{dev}}, nil
In pkg/cdi:
// UpdateAnnotations updates annotations with a plugin-specific CDI device
// injection request for the given devices. Upon any error a non-nil error
// is returned and annotations are left intact. By convention plugin should
// be in the format of "vendor.device-type".
func UpdateAnnotations (annotations map[string] string, plugin string,
deviceID string, devices []string) (map[string]string, error) {
In experimental code
        for , device := range response.CdiDevice {
            deviceID := strings.Replace(strings.Replace(device, "/", "-",
-1), "=", "-", -1)
            key, err := cdi.AnnotationKey(driverName, deviceID)
            if err != nil {
                return fmt.Errorf("could not get annotaion key, plugin: %s,
err: %+v", driverName, err)
            value, err := cdi.AnnotationValue(response.CdiDevice)
            if err != nil {
                return fmt.Errorf("could not get annotation value, devices:
%s", response.CdiDevice)
            }
            annotations = append(annotations, kubecontainer.Annotation{Name:
key, Value: value})
            klog.V(3).Infof("pod: %s, container: %s, annotations: %s",
pod.Name, container.Name, annotations)

    Q: Do subdirectories get parsed for CDI specs:
```

- Scanned here
- Seems to only scan single top-level folder
- Create a ticket to address this.

Jul 12, 2022

Attendees:

- Alexander Kanevskiy
- Ed Bartosh
- Ian Forbes

- PR for client configuration #67
- For review / discussion:
 - Spec extension: Add HostPath to Device spec to match how devices are handled in k8s:
 - Would ideally tag a release before this is merged so that 0.4.1 is available for consumption without requiring the spec update.
 - Spec extension: Add "meta" devices #71
- We have added CDI support to the NVIDIA Container Runtime for use in situations where a CDI-capable runtime cannot be used.
- PRs discussions -> we will move it to the next meeting when more people will be participating.
- Ian: interest in trying CDI/Dynamic resource allocation
 - Ed's and Patrick's branch:
 https://github.com/bart0sh/kubernetes/tree/PR0100-dynamic-resource-allocation

Jun 28, 2022

Attendees:

- Evan Lezar (NVIDIA)
- Alexander Kanevskiy (Intel)

Agenda:

- Meeting series
 - Message has been sent to CNCF (responsible person is on leave)
- PR for client configuration #67
 - o Is it ready for review?
 - Blocking rename of master to main
 - o Will ping Ed on Slack.
- Annotations as part of CDI spec #30
 - Zvonko Kaiser will present a use case in the issue and open a PR
 - These would be container-level annotations and not pod-level annotations
 - Should we only allow annotations with a specific key format i.e. not override known annotations like k8s annotations.
 - These would be CDI spec annotations.

Jun 14, 2022

Attendees:

- Ian Forbes (vmware)
- Jukka Rissanen (Intel)
- Evan Lezar (NVIDIA)
- Alexander Kanevskiy (Intel)

Agenda:

Introduction to COD and focus on CDI and KEP

- Ian Forbes from VMWare discussing their use case
 - Pointer to Google's presentation on KubeCon: https://sched.co/ytlt
 Improving GPU Utilization using Kubernetes Maulin Patel & Pradeep Venkatachalam, Google

0

- PR dependencies #68
 - Krisztian will approve and merge

May 31, 2022

Attendees:

- Eduard Bartosh
- Alexander Kanevskiy
- Krisztian Litkey
- Jukka Rissanen
- Evan Lezar
- Patrick Ohly
- Markus Lehtonen
- Kate Goldenring

Agenda:

- Status update on KEP and implementation
 - Patrick:
 - KEP review during KubeCon
 - Decision to go with custom object to hold user parameters: better validation
 - Potential ideas on autoscaler integration
 - \circ Ed
- Prototype implementation ongoing
 - Waiting state for Pods with resource references
 - Missing communication with plugin TODO
 - New gRPC service vs. adding methods to existing Device Plugins API gRPC proto.
- Not yet ready to be reviewed by kubelet maintainers
- PR for client configuration #67
- PR dependencies #68

May 17, 2022 - Meeting canceled due to KubeCon Europe

May 3, 2022

Attendees:

- Eduard Bartosh
- Krisztian Litkey
- Alexander Kanevskiy
- Evan Lezar

- [Ed] Kubernetes prototype, plugin API
- [kli] PR #56 Merging. Done
 - Merged
 - o Evan will need to update podman implementation in conjunction with with tag bump
 - Will tag once #64 is merged
- v0.4.1 tag
 - Will wait for #64 to merge
 - Will update readme for runtime configurations #67
 - Will have podman version bump ready with AutoRefresh before dependabot picks up tag update (Evan)
- Rename master to main
 - o Once open PRs are merged Elezar will change / rename default branch to main

Apr 19, 2022

Attendees:

- Evan Lezar
- Ed Bartosh

Agenda:

- Containerd release:
 - Should we backport CDI support to 1.6?
 - o https://cloud-native.slack.com/archives/C4RJZ9Z6Y/p1649895908929599
 - Ed to add configuration / documentation for Containerd to address issue comments:
 - **#**59
- Session recordings:
 - Podman demo from Feb 8, 2022
- Spec extension: #65 (see #45)
 - Add a "Type" field to allow for tmpfs mounts, for example
 - What is required to write out a new spec version?
- v0.3.2 tag released to address code size growth in podman
 - See https://github.com/containers/podman/pull/13852
- Evan Lezar DE comment on #56
 - Once merged, tag and update podman implementation to use AutoRefresh=false

Apr 5, 2022

Attendees:

- Evan Lezar
- Ed Bartosh
- Patrick Ohly
- Alexander Kanevskiy
- Jukka Rissanen
- Kevin Klues
- Mike Brown

- PR status updates
 - Merge after test confirmation
 - Should we wait on Mike Brown's ok?
- Discuss spec & API versioning policies to address this review comment. Create requested github issues.
 - We do already validate the version of the spec in the API
 - We will create an issue to discuss spec upgrades across versions
 - We will include tests for loading old specs
 - We will provide functionality in the API to write out specs (PR incoming)
 - We will only write out specs in the current version
 - Evan Lezar DE | will create an issue and also update the containerd PR comment linked above.
- Add config options for different runtimes:
 - Evan Lezar DE Podman
 - o Cri-o
 - Containerd
 - Ed Bartosh Created container-device-interface#59

Mar 22, 2022

WARNING: Due to US move to summer time, and meeting invite was from US timezone calendar, the schedule is screwed today. We will start at the usual "european" time slot.

Attendees:

- Krisztian Litkey
- Evan Lezar
- Alexander Kanevskiy
- Jukka Rissanen
- Patrick Ohly
- Kevin Klues

_

- Should we update the meeting timezone to be "local" i.e. Finland?
 - Agreed: move to "EU" timezone calendar, so the meeting will be following the European daylight saving time schedule. Sasha to send an update to CNCF.
- PR status updates
 - o CRI-O:
 - Merged.
 - Needed update to new tagged version (crash fix missing).
 - Will wait until Thursday to see what containerd says about PR.
 - Containerd
 - What about the versioning question?
 - elezar will comment on the versioning question and ask if there are specific requirements on versioning of the API given the responses from kli, ed, elezar. comment

- Tagging current state:
 - Evan to tag 0.3.1 with current "crash fix": Done v0.3.1
 - If containerd ok with "self-refresh", we will tag 0.3.2
- PR#56 / refresh/locking:
 - containerd/CDI feedback.
 - Default behavior is to self-refresh (as before)
- Kevin Dynamic resource allocation KEP
 - Adding support for multiple plugin versions to device manager
 - v1/beta1, v1/beta2
 - Get rid of empty messages. Allows information to be passed in both directions. Allows for extension and deprecations?
 - Target for PR is end of week
 - Question: Can we generalize plugins in the kubelet? (CSI, device-plugin)
 - Abstract out "talking" to the plugin and reuse passing things to the device manager. This may be applicable to multiple plugins
 - Ouestion: Can we add extra fields?
 - This will be step 2 of adding v1/beta2. Adding fields doesn't require a version bump in this case.
 - o PR out but not merged before release (Apr 19, 2022)
 - Should not require a KEP
 - Is there a plugin that gets run in CI:
 - There is a dummy plugin that is run in CI that goes through the lifecycle of the plugin.
- Patrick Changes in branch
 - o API Changes are done
 - Resource driver with immediate allocation
 - o Resource claim controller done
 - Working on Authz
 - Next: Scheduler plugin

Mar 8, 2022

Meeting is canceled: some people are on vacation, no agenda.

Feb 22, 2022

Attendees:

- Ed Bartosh
- Evan Lezar
- Jukka Rissanen
- Alexander Kanevskiy
- Markus Lehtonen
- Krisztian Litkey

- Slack communication:
 - We just use <u>CNCF #tag-runtime</u>

- Do we want to create a channel elsewhere? Where would it make sense?
- Tagging the <u>v0.3.0</u> release (Evan Lezar DE) will tag the version after this meeting)
 - Still pending podman update (go mod reference updated to v0.3.0): https://github.com/containers/podman/pull/13317
 - o cri-o can be updated at a later stage (since the CI is green now)
- Status updates of PRs (containerd/cri-o)
 - o cri-o:
 - PR is approved
 - waiting on CI fixes to land before merging PR
 - Krisztian will ping Peter to ask about merging before CI changes are in
 - We will give the CI changes until next week to land before trying to get the changes merged without them.
 - containerd:
 - changes are in place (adapted from cri-o): https://github.com/bart0sh/containerd/commits/cdi5
 - waiting for cri-o changes to be merged
 - <u>issue</u> running k8s from latest release (currently based on 1.5.7)
 - Let's not create confusion by mentioning NRI. If NRI is raised we should make it clear that this COULD be implemented as an NRI plugin, but there is no clarity on the timeline of when that will land.
 - Add PR to agenda of Containerd meeting on Mar 10, 2022 and discuss the PR there.
- CDI Device Plugin
 - Would it make sense to implement this in the kubelet instead of requiring a plugin?
 - At this stage being decoupled from the kubelet release cycle is important?
 - We could consider adding a separate repo in the GitHub org (container-orchestrated-devices)

Feb 8, 2022

Attendees:

- Evan Lezar
- Alexander Kanevskiy
- Patrick Ohly
- Markus Lehtonen
- Jukka Rissanen
- Ed Bartosh
- Krisztian Litkey

- Podman CDI demo (elezar)
- Tagging version of CDI (see previous discussion)
 - o SPEC.md
 - Preamble
 - Example
 - Release tracking (is this relevant)
 - Plan of action:
 - Remove unused fields as a cleanup (e.g. runtime) Evan Lezar DE

- Update references in source to 0.3.0 and tag as 0.3.0 Evan Lezar DE
 - #47
- Update podman to point to 0.3.0 (currently v0.0.0-20220111162300-46367ec063fd) Evan Lezar DE
- Man page for spec [container-device-interface#44]
 - Request from cri-o documentation folks for a manual page
 - Might be requested by the containerd maintainers too
- CI for testing
 - Test should run on every PR
 - o Ed will create an issue (can review / discuss on Feb 22, 2022)

Jan 25, 2022

Attendees:

- Eduard Bartosh
- Kevin Klues
- Alexander Kanevskiy
- Krisztian Litkey

Agenda:

- PR Review:
 - https://github.com/container-orchestrated-devices/container-device-interface/pull/43
 - Evan or Kevin would look at it
- [ed]: Deactivate call in device plugins API
 - Still applicable to Device Plugin APIs
- Runtime CDI support
 - Krisztian has draft for CRI-O
 - Ed working on containerd
- Patrick: KEP status update for Dynamic Resource
 - Autoscaler concerns and handling those.

Jan 11, 2022

Attendees:

- Alexander Kanevskiy
- Evan Lezar
- Jukka Rissanen
- Krisztian Litkey

- API Cleanup PR: <u>container-orchestrated-devices/container-device-interface#39</u>
- Tagging current version
 - Tagging is cheap, so it may make sense to tag the current version to make it clear that this is being used by e.g. podman.
 - Suggestion: Use minor version number to indicate spec version (data model) and update patch number to indicate an API / functional change. Consumers can then select a

- specific spec AND functionality. If there are behavioral changes then due to the spec, we bump the minor version.
- Tagging preparation: we need to make sure that SPEC.md and example reference the current version number.

Notes:

- Containerd community meeting on Jan 13, 2022 (Evan / Kevin to attend)
 - We should have confirmation as to whether an NRI plugin-based implementation will be used going forward.
 - o If NRI is not used, the existing CDI PRs by Ed would have to be updated / pushed.
- Question: Is the plan to also add an NRI implementation for podman?
 - o NRI is not really targeted at serverless runtimes and as such may not be applicable
 - The plan for the time being is then to use CDI directly in podman and use NRI for other runtimes (containerd, cri-o)

Action Items:

- Alexander: Upload recording form 14 December 2021
- Elezar: Prepare podman demo for next meeting
 - At least with existing API
 - o Ideally with PRs for new API
- Elezar: Make a suggestion for tags to apply for spec / API versions

Dec 28, 2021

Meeting is canceled: most of attendees are on vacation.

Dec 14, 2021

Attendees:

- Alexander Kanevskiy
- Eduard Bartosh
- Kevin Klues
- Jukka Rissanen

Agenda:

- Patrik/Kevin' talk proposal is not yet submitted, but should be done soon.
- Ed/Krisztian: prototype of CDI and NRI plugins
 - o Simple device plugin that discovers all CDI json's and announces that to kubelet
 - Plugin injecting device via container annotation (return of Allocate())
 - Two implementations:
 - Containerd quick hack-in by Ed
 - NRI plugin by Krisztian

Nov 30, 2021

Attendees:

- Alexander Kanevskiy
- Eduard Bartosh
- Evan Lezar

- Patrick Ohly
- Krisztian Litkey
- Jukka Rissanen
- Kevin Klues
- Markus Lehtonen

- Discuss issue #38 / PR#39 improvements of CDI APIs
 - o Agreement to drop support of short names for now. Might reconsider later.
- Can we set the (fully qualified) device id in a k8s device plugin and have the required changes reflected in the OCI spec:
 - Ed / Krisztian will coordinate on setting up a test case with a device plugin that makes these modifications and triggers a CDI-enabled runtime (possibly as an NRI plugin) that resolves these IDs and makes the required OCI modifications.
 - Either modify the device host path or annotations?
- Dynamic Resource Allocation KEP: ready to be proposed to SIG Node

Nov 16, 2021

Attendees:

- Alexander Kanevskiy
- Eduard Bartosh
- Krisztian Litkey
- Markus Lehtonen
- Evan Lezar
- Patrick Ohly
- Kevin Klues

- Approval / merging of Pull Requests
 - No one merges their own code
 - One must not merge PRs from your own org (reviewers should be selected from other orgs)
 - o If there are no changes requested the (last) reviewer approves and merges the PR
 - Use #tag-runtime slack channel (https://cloud-native.slack.com/archives/CPBE97SMU)
 for discussion
 - Evan Lezar DE update
 https://github.com/container-orchestrated-devices/container-device-interface/blob/master/
 /CONTRIBUTING.md with the contents above
- Caching of parsing CDI json files
 - The spec should not define when the spec files are loaded. The API should be updated to only operate on the spec and the lifecycle (triggering reloads etc) should be left to the consumer.
 - Loading / monitoring functions can be provided as utility functions but these are orthogonal to the spec.
 - The CDI root should be configurable (list of directories) with a fallback to the default (/etc/cdi)
 - We should focus on the fully-qualified device ID (vendor.com/type=id) implementations instead of allowing short IDs (just "id") to be used

- o The meta devices (e.g. "all") can be implemented as its own spec
- Question: There does not seem to be a way to go from a single device ID to a specific file that needs to be loaded.
 - This is by design as the spec depends on different parts which could be generated with different life cycles (static vs dynamic)
 - Let's work on an implementation first and then optimize the implementation if required.
- Containerd integration
 - o PR is out: https://github.com/containerd/containerd/pull/6204
 - We need reviews from containerd
- Kevin / Evan to look at NRI
 - NRI PR set
 - extended NRI
 - <u>CRI-O integration</u>
 - containerd integration (not in sync yet with the other two)

Nov 2, 2021

Attendees:

- Ed Bartosh
- Alexander Kanevskiy
- Mike Brown
- Krisztian Litkey
- Markus Lehtonen
- Antti Kervinen

Agenda:

- [ed]: CDI and way to go: into core or as NRI?
 - Mike: integrating to tasks will conflict with earlier design decisions for different clients of containerd.

Oct 19, 2021

Attendees:

- Alexander Kanevskiy
- Kevin Klues
- Evan Lezar

Notes:

- We have spec and library and podman implementation
- Containerd and cri-o integration in progress
- Working on KEP for new device plugin
- NRI
 - o Implementation in progress
 - Cri-o is working, but containerd is still in progress.
 - Needs further working / refactoring of implementation in containerd

- Idea/demo: CDI implemented as NRI plugin to test interface (although having it as core functionality makes more sense. The reasoning is that plugins are more difficult to debug especially in the context of devices).
- o Question: Has there been an attempt to reimplement static CPUManager in NRI
 - CRI Resource Manager has been demonstrated that can act as a drop-in replacement with a static policy. POC has been implemented as NRI plugin
- o Links:
 - https://github.com/containerd/nri/pull/16
 - https://github.com/containerd/containerd/pull/6019
 - https://github.com/cri-o/cri-o/pull/5318
- Logos
 - Evan to follow up with Renaud on the original files

- Welcome to restarted series of meetings.
 - Every two weeks until June 16 2022
 - Recordings are automatic and will be uploaded
- Status update on our activities
- Logo?

June 29, 2021

Notes:

- Meetings for July are cancelled (vacations)
- Feedback on KEP proposal
 - @elezar to have a look at the PR
- Feedback on CDI Demo (reusing CSI primitives)
- Device permissions in CDI (ed)
- Who will the NVIDIA representative be (elezar)

June 15, 2021:

- [kli] Prototyping NRI with an Extended Scope
 - Related <u>KubeCon Talk</u>

May 18, 2021:

- [pohly] Kubernetes API and scheduler support for CDI:
 KEP draft in https://github.com/pohly/enhancements/pull/1
- [pohly] Present that work to the wider community at KubeCon NA how?

May 4, 2021: Cancelled in favor of KubeCon

March 9th, 2021

• https://github.com/cncf/sig-runtime/issues/24 Any action items for this?

Feb 23, 2021

- Logo for CDI
- CDI First implementation:

https://github.com/container-orchestrated-devices/container-device-interface/pull/19

@Kad/@ed/Ukri wants to show a demo!

Code: https://github.com/intel/proof-of-concept-cdi

Mikko: devices for non-root containers: https://github.com/kubernetes/kubernetes/issues/92211
 Proposed patch:

https://github.com/mythi/containerd/commit/0918c1aa3096e0f16b0e3fd3abed9f5554a214df

Feb 9, 2021

Renaud: logos

kad/ed/ukri: demo of our PoC

- Podman implementation discussion:
 - API for podman
 - o Loading the CDI files and keeping them in memory (global vars, passing a context, ...)
 - Spec golang representation (e.g: Do we want to redefine the OCI structures)
 - Applying changes to the OCI spec

•

Non-root containers and devices:

https://groups.google.com/g/kubernetes-sig-node/c/zhOxWdGyY3Qhttps://github.com/kubernetes/kubernetes/issues/92211

Jan 12, 2021

- Wins of 2020
 - CDI first G.A draft
 - Kubecon presentation
 - Pod Resources API
 - Resource Management Wg
- NRI and CDI
 - Mike: Let's start to get up on podman, need to keep the links between the two projects so that the injection points are at roughly the same place
 - Urvashi: Agrees with that approach
 - Kad: Similar line of thinking
 - Let's do the fast path as a proof of concept
 - NRI is still missing a few items

- Use Cases:
 - NVIDIA GPUs and Mellanox NICs, Intel FPGA and GPUs
- Kad: Make sure that NRI is
 - a powerful : https://github.com/containerd/nri/issues/3
 - performant interface time budget: <1s, e.g. as described cdi is inside NRI https://github.com/containerd/nri/issues/2 to avoid much of "fork/exec" calls then scaling.
- Discussion on exposing entities to upper layers
 - o Container refers to CRD and that reference is kept at the container runtime level

December 07, 2020

- Nikolay Nikolaev: service mesh and accelerator devices, CDI questions.
 - Walked through last time
- CDI Next Steps
 - Is it podman and containerd through NRI?
 - Is it podman with bare CDI integration and containerd with NRI integration
 - Think the question comes in terms of CRIO and Containerd
 - Do we want a flag to enable it, with podman this is easy
 - Regarding kubelet it's possible to do several approaches
 - We can reuse the device plugin functionality for now
 - Need a new proposal for something better:
 - Existing device plugins pass the absolute path
 - We could distinguish with absolute path vs non existing path
 - Passing down in the CRI a device with name "vendor.com/device=myDevice"
 - Al: Making a presentation + starting a rough document that leads to a kep
 - Have a proposal && proof of concept
 - o Concretely what are the next steps
 - Podman: Need to get a PR up and running
 - Containerd: NRI
 - Containerd -> NRI -> CDI
 - NRI -> CDI
 - Comments about NRI
 - Ping Michael Crosby: https://github.com/containerd/nri/issues/3
- Kubecon Europe: Do we want to submit something?
 - We are less likely to get a spot given we just presented in US the same topic
 - o Do we have significant progress or new ideas we want to present?
 - Deadline is next Monday && it's a lot of work
 - o Panel or Talk?
 - PIC: Alexander Kavnevskiy
 - Keep the group up to date on cncf-sig-runtime
- Move next meeting to Jan 12th

November 23, 2020

- Post KubeCon, not that many attendees
- Nikolay Nikolaev: service mesh and accelerator devices, CDI questions.

November 10th, 2020

- https://github.com/containerd/nri
- [KAD]Try to replicate the hooks in CRI-proxy in CRI implementation
 - Exec style is expensive
 - We need a hook in the create path
 - Chain of plugin scenario (CDI + CPU resource plugin)
 - o Update calls in the CRI, e.g. kubelet asks CRI to Pin to CPU 1,2,3, CRI pins to CPU 22
 - Should we change kubelet to ask to pin an amount of CPU rather than specific CPUs
 - o Communication channel between kubelet and CRI implementation
 - o Ability to order the Update message and read the container state
 - Plugin should be able to alter other existing containers in the Create/Start/Update path
 - You have a container with X allocated CPUs and you free the CPUs
 - o Al: Open an issue
 - Added <u>issue</u> explaining what we're after
- KEP draft: Moved to next meeting

October 12, 2020

- Next concrete steps for CDI and NRI
 - What level of degree do we want to extend the NRI for the intel use case, hooks, different stages for pods?
 - Initial design in NRI is around resource management (CPU and NUMA) and expanding out to huge page support and cache (over the past month)
 - o Lifecycle hooks (create, start, delete and update) have worked out well
 - No current Pre-create but can see use case to modify the spec + configurations
 - Mike: Currently be flexible enough today
- Is the NRI something we will see in other runtimes?
 - Mrunal: Need to understand the bigger picture
 - When are we calling into the NRI?
 - o Michael: Lots of low level details shouldn't be in the kubelet
 - Alexander: React on pods, individual containers, trigger updates of CPU cores
 - o Mrunal: If it's over the CRI it's easy to integrate
 - Alexander: Modify and intercept messages (e.g. CreateContainer)
 - NRI: Can support devices and resources and anything we can think of in the future
 - Alexander: Want a gRPC (or other) where we integrate closer to the runtimes for high performance invocation.
 - Mikeb: Store what we receive from Kubelet before modifications are made, and maybe after modifications are made
 - Alexander: Get from the runtime the whole state of the system

- Michael: Want one call in the kubelet system to feed us the whole pod specification to be able to make more efficient decisions
- Mikeb: What about the other objects related to the spec? Lots of stuff before the calls are made (caching, pulling, ...). Kubelet has a "cm" (container runtime manager) that has a lot of information the CRI implementation might need.
- Mrunal: Came up in the sidecar conversation (startup and shutdown order), might help make a case for the podSpec CRI call.
- Alexander: Need to define responsibilities of the Kubelet and CRI implementation.
 - Kubelet says what, CRI implementation says how

Containerd

- o CRI shim stays the same so that it can provide the pod information
- Need to update the containerd client in docker with NRI
- Could support this at the ctr level

Podman

- Need to look into NRI, might need a demo
- Next concrete steps:
 - Look into writing an NRI CDI plugin
 - Look into NRI with docker
- Long term CDI will support pod level devices

Kubecon panel discussion

- o Format: Q&A between speakers, little slides, keep 10-15 minutes for external questions
- Need a moderator to setup the tone: Renaud
- List of Questions:
 - Introduction! Renaud + everyone introduces themselves
 - What? Alex and Renaud Representatives from different companies and areas, what do you expect from, what problems are you trying to solve?
 - From CDI and from the COD workgroup
 - How different? Urvashi Introduction to CDI, mental model, how does it look like from a kubernetes end user perspective
 - How it's all connected? Where are we involved? Scope? Mike and Mrunal How does the NRI, runtime hooks and CDI ideas intersect?
 - Roadmap? Renaud as a followup of setting the tone What our roadmap for the COD workgroup
 - Before we go to audience questions
- What are we getting out this is a slide
 - More Involvement
 - Share the ideas, get some feedback
 - People from HPC, smaller runtimes, other sigs, non kubernetes users, "strange devices" people
 - Intersection with other resources and
- Present why this important at first
- Setup the doodle ASAP

С

- Next meeting VM based runtimes
- [Kad] Antti started drafting the kep about CRI-kubelet communication
 - https://github.com/askervin/resource-management-improvements-wg/tree/58r_kep_resource_to_cri/keps/sig-node/NNNN-pass-resource-reqs-limits-to-cri
 - Ai: Formal review next meeting

September 29, 2020

- Next concrete steps for CDI and NRI
- [Kad] Antti started drafting the kep about CRI-kubelet communication
 - https://github.com/askervin/resource-management-improvements-wg/tree/58r_kep_resources to_cri/keps/sig-node/NNNN-pass-resource-regs-limits-to-cri
- [Renaud] have a draft podman example for CDI
 - https://github.com/containers/podman/compare/master...RenaudWasTaken:cdi-draft?exp and=1
- [Mike] Would be nice to have a Generic plugin system in podman / containerd, maybe through the NRI or an additional mechanism
 - o [Kad] NRI is specified and scoped to OCI hooks and containers, may need more
- [Mike] What is the Pattern/Architecture for devices at the pod level, what do the changes in kubernetes look like?
 - [Renaud] No changes on the pod spec side, users will still request devices
 - [Renaud] For pod level devices, that would require a change since this is not currently possible
 - [Renaud] I mostly expect us to change kubelet to pass the device information in the CRI message
 - o [Kad] We could use the current device plugin interface to pass that information down
 - [Renaud] Wondering if it makes sense to remove the ability for the device plugin to change the specification
 - [Kad] Has a use case where it makes sense to change the pod/container spec from the device plugin.
- Mrunal and Mike are working on a kep to G.A the CRI
 - Minor Extensions / adds would not be an issue for V1
 - Cgroups v1/v2 specifics in CRI APIs
 - Coexistence of V1 and V2? Should it be hidden to the kubelet?

September 15, 2020

- Alex discussion around new repo for discussions around SIG-NODE interactivity with Node Resources, higher level scheduling and state abstractions needed for qos scenarios. URL: https://github.com/container-orchestrated-devices/resource-management-improvements-wg
 - Alex discussed the status of CRI api (currently alpha2, possibility of adding additional capabilities for CDI/NRI scenarios regarding additional node level resource requirements of sandboxes/container prior to CRI beta.
 - Will work with appropriate SIG*s to ensure teams come together for these higher level sandbox scenarios.
 - AT&T example discussed: how to manage system resources at the kubelet level (cadvisor++) vs lower level admin oriented management (allocate certain resources to particular processes container runtime / kubelet and other node resources) (of course not just cpu and memory)

0

• Admin decision, group messages will be done on cncf sig-runtime: url:

• Will invite Michael Crosby for the next call to discuss next steps, status,.. How best to arrange the various workgroups.

August 2nd, 2020

- Node Resource Interface
 - https://github.com/containerd/containerd/pull/4411
 - https://www.youtube.com/watch?v=MVIBXUafB38
 - NRI is trying to solve QoS
 - Hardcoded in Kubelet
 - Similar interface to CNI, managing Node Resources like networking
 - Hooked into CRI / at the container runtime level
 - Can be expanded for generic device management
 - Need more of a way to hook in the lifecycle of a container
 - Exposing to different lifecycle hooks, updates, resizing of a pod
 - When we deal with the runtime spec, is it mutable or immutable?
 - Mrunal: It is mutable today for Update calls but would prefer we keep it immutable after writing config.ison.
 - Mike: Mutable pre-create, immutable post create
 - Update Resources in CRI might require changing the spec
 - Need to propagate these changes back (above CRI)
 - CRI has an in memory representation, containerd/podman builds the memory representation, CDI/NRI edits are applied, config.json is written... Mike: Edits may need to be validated or some edits rejected for security.
 - Spec is returned, probably unused today (might require structuring?)
 - MCrosby: The more you do in Docker the less you can do at the lower level (which is typically more robust)
 - Need an error flow so that we can reject a pod rather than have a Pod Runtime Failure
 - Mrunal: Other features requires this too
 - Pod Scaling & Deployment
 - Two step scheduling, scheduler tells us what nodes to run on
 - At the node level, we decide best placement
 - Possible use case for improving the scheduler
 - Next Steps
 - Hooks that allow you to modify the spec before create? Pre-create
 - When the kubelet asks to update the cgroups (update CRI call). Need an UpdateJSON call in the CRI?
 - How can we merge these two Specs? https://github.com/containerd/NRI
 - Decoupled from containerd where we can collaborate
 - Easy import
 - MyContainerRuntime run --device MySuperDevice
 - Based on the CNI model, is there some intersection with CNI?
 - CNI is typically based on pids and network namespace
 - NRI typically has access to much more information
 - Need to think through the interface much more if we want to think through invoking a CNI plugin (or a compat layer)

- NOTE(AdrianC) CNI performs operation at Pod level not container level, it needs to be invoked once per pod. Some CNIs rely on pod uid to cache data for cleaning when the Pod is deleted.
- First class chaining
 - See or get information that the previous plugin sets
 - Use case: Customer specific labels for QoS into confined labels (exclusive CPUs, hyper threads, ...)
- Could we move security policies from kubelet into the container runtime
- Conflict between create and Update
 - E.g: Create -> adds a device node (kubelet not aware)

CDI Roadmap

- Note that these slides will be presented on Tuesday at 8am PT at the <u>CNCF TOB</u> (feel free to edit before that!)
- o Defining Alpha / Beta / G.A general guidelines
- Engaging with other ideas (NRI, Network Storage)
- Writing the first Demos
- Other Runtimes (LXC, HPC runtimes, ...), presenting at the HPC Containers Advisory Council on Thursday

Pull Requests

- o Common Vocabulary Terms
- Golang Specs
- o JSON schema

July 21st, 2020

- Hooks pointing to only executables
 - This is the case in runc to prevent security issues?
 - Issues in the past where someone can escape?
 - For StartContainer it seems like you might need to follow symlinks.
 - Call stat on the file, windows/unix how to identify executables?
 - o Need more input and research?
 - Should "runtime" hooks be in a root owned directory (e.g: /etc/cdi/bin)?
 - Concern about it being a "trashcan"
 - Vendors might want to install binaries in /usr/local/bin
 - We can also start with just executable and then expand
- Allow Device Nodes or Files to be mounted from other containers
 - Device node wouldn't work
 - Mounting files from other containers is a work in progress (CSI plugin that does it today)
 - CSI code is responsible for setting the container and CRI does the bindmount
 - CSI sets up the rootfs
 - Think of this more as an Image mount rather than a container mount
 - A volume of type "image", CSI is responsible for pulling the image
 - @Mrunal @Kir to share a presentation when that work is ready

Ordering of mounts

- Order mount operations logically, directory-wise
- We don't have a use case for masking a directory (hiding a previous mount by putting a new mount on top of it and the new mount is shorter)

Redirecting output of hooks

- What is the use case? Debug?
- If the hooks are executed by CRIO or containerd, hooks will be executed by runc and we don't have the flexibility in the OCI spec.
- Would need to open an issue on the OCI spec
- Today hooks redirect themselves stderr and stdout
- o Redirecting the output to some file on the host is going to be host local
- Decided not to address this for now

Multiple resources for a specific vendor

- Current spec points to a single file per vendor, should this file represent a vendor or a resource?
- Need to allocate a VF and pass it to the container
- o differentiating between network resources because they are simply wired differently
- Need to pass different mounts (e.g: RDMA enabled NIC)
- Create per resource JSON files, on the one hand kubelet will allocate specific device ID and the runtime will mount whatever it is in the device spec
- Multiple resources per vendor (same scenario for FPGAs), each FPGA function will have its own separate file
- Do we allow injecting environment variables?
 - If we mount multiple resources to a container how to know which mount is owned by which resource.
 - It being an actual capacity of the device plugin is a strong argument for it to be part of CDI

Node Resource Interface

- Will be presented at sig-node at 10pm
- Similar to how CNI works today
- Concerns about dynamic updates of resources
- o Github Pull request is here
- o "Runc update" of resources kinds of break the model
 - Requires you to re-implement your changes
 - Process will likely see resources be changed
- Any way to mimic StartContainer (e.g. Idconfig)?
- Pivot Root is getting executed as part of the create operation, which makes it difficult to add new mounts (this is specified in the OCI spec).
- @Kad will be in holidays for the next three weeks

June 23nd, 2020

Attendance/Updates/Standup (PLEASE ADD YOURSELF):

- Adrian Chiris
- Marek Counts
- Mike Brown (IBM)
- Renaud Gaubert
- Alexander Kanevsky
- Ed Bartosh
- Mrunal Patel

Agenda:

• [Renaud] Zoom + Recording

- Zoom is setup until December
- Need to figure out recording
- [Renaud] Charter of the group
 - Needs to finished by Friday
 - Need to have a separate use case section in the Charter Doc
 - Figuring out the different use cases should be part of the charter?
- [Renaud] Kubecon USA Panel
 - o Deadline is Sunday, June 28, 2020
 - Needs to also be finished by Friday
- [kad] <u>Usage of the devices by non-root containers</u>.
 - Can we copy the user ID of the process to the user ID of the devices?
 - CRI is missing the "hostPath Devices permissions"
 - The expected device was for kubelet to create the device node?
 - Special devices with very specific privileges?
 - o Do we have edge cases where we would be breaking users?
 - O What are short mid term solutions?
 - Annotations are passed down to containerd/CRI-O that indicates whether to use the user process permissions for devices?
 - Start with the annotation? Optin flags that we can set as we want. Annotations can be passed through the device plugin.
 - What are the scenarios where the gid/uid are relevant in containers? Legacy workloads?
 - Discuss whether we need a kep or not? We need keep sig-node informed on why we are doing this.
- [Renaud] Monitoring of Devices in Kubernetes
 - No specific topic, just want to bring this up as we are moving towards G.A for this topic in Kubernetes, might be interesting to have some feedback on missing use cases
 - o Disable Device Metrics in Kubelet
 - G.A Plan for Monitoring Compute Devices (in K8s)
 - GPU_UTILIZATION{gpu=UUID} 100
 - With Pod Resources: GPUI_UTILIZATION{gpu=UUID, container=Foo, pod=Bar} 100
 - Ukri: Don't have the level of detail, but would like to have it
- [Renaud] CDI
 - Github Org has been created
 - Let me know if you aren't in the github organization
 - What are the next steps?
 - We need to re-use the spec as much as possible
 - Clone the CDI base / Make a Pull Request
 - List of items which are allowed to transform the spec, shouldn't be able to transform all of OCI?
 - Devices, Mounts, Hooks,
 - Runtime should not blindly process all the items
 - We should reuse the fields that are existing from other runtimes
 - Try to reuse the golang structs
 - Ordering needs to be specified, this is pretty important for the mounts
 - Deliverables:

- A specification
- A golang API
- A golang library for transforming the OCI spec using the CDI spec
- A JSON schema for validating JSON specs
- CLI examples
- Simple jq might work?
- o Lifecycle of the project?
 - Annotations as a way for enabling / disabling this feature
 - Multiple Phases: Draft specifications, POCs & validation with existing devices
 - Feedback from the kubernetes groups (Rancher/Openshift/...) + the class of users not using CRI
 - Parallel to CNI
- o Is this scoped to Filesystem resources?

June 10th, 2020

Attendance/Updates/Standup (PLEASE ADD YOURSELF):

- Mike Brown (IBM)
- Adrian Chiris
- Moshe levi
- Alexander Kanevskiy
- Ed Bartosh
- Ukri Niemimuukko

Agenda:

AI:

- @Kad: Doodle Request
- @Renaud: Kubecon Google Doc
- @Renaud: Creating a Github organization (COD)
- [Renaud] Administrative stuff:
 - Meeting day, time and recurrence, time zones
 - Once every other week
 - Google Groups
 - Recording and Zoom meeting I have covered
- [Renaud] Charter of the group (getting approval)
 - Are we aligned on short term and long term goals?
 - o Do we have a list of concrete deliverables for this group?
 - Discussing with the Network Plumbing Workgroup?
 - How to pass device information to the container?
 - Provide a mapping + additional info between the device and the network it is associated with
 - Ping pong between device plugin and the Multus CNI plugin / Meta-CNI
 - Downards API vs Fixed mount point?

- https://docs.google.com/document/d/1rBm-L1ymXIjoKNA6w2lixwBAjt9-tnbs-ee
 7AcLRbVE/edit
- Discussing with Sig-Storage?
 - Might have same use cases
- Kad: Improvement of devices in Kubernetes
- o Adrian: Do we have a list specific use cases to help?
- https://docs.google.com/document/d/1wPIJL8DsVpHnbVbTaad35ILB-jqoMLkGFLnQp WWNduc/edit#
- https://docs.google.com/document/d/1Tc0Kc4GDWx1gFvGQbBUizudSuND6Kq8GiH7 KVm X5eg/edit#heading=h.5lakm98lya8j
- [Renaud] Kubecon USA
 - o Deadline is Sunday, June 28, 2020
 - Submission of a panel
- [Renaud] CDI
 - Administrative
 - Create a Github (COD Github)
 - Review process? Probably pull request
 - Deliverables:
 - Is it a specification or not?
 - Resources -> Kubernetes?
 - Specification -> OCI?
 - An example could be: https://github.com/opencontainers/runtime-spec
 - It has the spec
 - It has a golang API
 - Bring this question with mrunal Patel, do they want governance over it?
 - Mike: It would be nice to have governance
 - Deliverables:
 - A specification
 - A golang API
 - A golang library for transforming the OCI spec using the CDI spec
 - A JSON schema for validating JSON specs
 - CLI examples
 - Notes about the CDI presentation
 - Containerd Namespaces
 - Can we use a specific field to detail what namespace this would be applicable to?
 - Should use the Kubernetes namespaces?
 - Common injection schema/system with (same schema as CRI-O)
 - Architecture questions
 - Where does it fit? At the runc/crun/kata level? At the docker/podman level? At containerd/cri-o level?
 - Runc has one source of truth, the OCI spec we want to keep it this way
 - Similar to the question How do you handle volume mounts inside a container?
- [kad] (maybe not for first meeting) Usage of the devices by non-root containers.

- $\circ\quad$ The Device Plugin interface was implemented where you expect to run as root user
- In kubernetes if you run a container as user e.g: 1000, it will mount the device with permissions that won't allow you to read/write from it