Minimizing Privilege for Building HPC Containers

Reid Priedhorsky

reidpr@lanl.gov

High Performance Computing Division
Los Alamos National Laboratory
Los Alamos, NM, USA

Timothy Randles

trandles@lanl.gov

High Performance Computing Division
Los Alamos National Laboratory
Los Alamos, NM, USA

R. Shane Canon

scanon@lbl.gov

National Energy Research Scientific Computing Center
Lawrence Berkeley National Laboratory
Berkeley, CA, USA

Andrew J. Younge

ajyounge@sandia.gov

Center for Computing Research
Sandia National Laboratories
Albuquerque, NM, USA

ABSTRACT

HPC centers face increasing demand for software flexibility, and there is growing consensus that Linux containers are a promising solution. However, existing container build solutions require root privileges and cannot be used directly on HPC resources. This limitation is compounded as supercomputer diversity expands and HPC architectures become more dissimilar from commodity computing resources. Our analysis suggests this problem can best be solved with low-privilege containers. We detail relevant Linux kernel features, propose a new taxonomy of container privilege, and compare two open-source implementations: mostly-unprivileged rootless Podman and fully-unprivileged Charliecloud. We demonstrate that low-privilege container build on HPC resources works now and will continue to improve, giving normal users a better workflow to securely and correctly build containers. Minimizing privilege in this way can improve HPC user and developer productivity as well as reduce support workload for exascale applications.

ACM Reference Format:

Reid Priedhorsky, R. Shane Canon, Timothy Randles, and Andrew J. Younge. 2021. Minimizing Privilege for Building HPC Containers. In *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '21)*, November 14–19, 2021, St. Louis, MO, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3458817.3476187>

1 INTRODUCTION

Linux containers have become a popular approach to develop, test, and deploy applications, because they let an application be packaged with its complete supporting software stack as a single unit, even if that stack is an OS distribution (the only exception being the kernel itself) [8, 26]. The simple foundation is that containers are processes (or groups of cooperating processes) with their own independent view of certain kernel resources, most importantly the filesystem tree.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SC '21, November 14–19, 2021, St. Louis, MO, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8442-1/21/11.

<https://doi.org/10.1145/3458817.3476187>

Recent challenges for large-scale scientific applications in HPC include the need for greater flexibility in supporting software stacks as well as application teams' desire to control the timing of changes to their stacks. Containers are an increasingly popular way to meet these needs,¹ and these benefits are not offset by a performance penalty. Studies of container performance impact have consistently shown that, when done correctly, containers introduce little to no performance penalty and sometimes even improve aspects like startup times [4, 16, 37, 40, 52, 54, 60].

While containers have been deployed across a wide array of supercomputers in recent years [7, 20, 36, 54, 61], containerized workloads still represent a distinct minority of supercomputer cycles [5]. We believe a key reason is the difficulty of building container images with unprivileged tools.

A long-standing principle for HPC centers across all security levels is to give users normal, unprivileged accounts, due to HPC's need for shared filesystems and their UNIX permissions. This "least-privilege" rule protects data from unauthorized access and the system from unauthorized modification, even if by mistake. In contrast, computing for industry web applications (a.k.a. "the cloud") uses single-user virtual machine sandboxes without shared filesystems [43, 58]. Data are protected by identity management, encryption, formal methods, and/or role-based authentication [10, 21, 22]. Such isolation techniques make it secure to allow container runtimes and build tools with elevated privileges. A conflict arises when these tools are proposed for HPC resources, where elevated privilege is almost universally prohibited.

This paper explores container build privilege in HPC, discussing how containers work, different container implementations' privilege models and how they relate to Linux namespaces, and why *running* unprivileged containers is easier than *building* them. Next, we analyze potential approaches to enable container build on HPC and compare related work in this area. We then detail two container build solutions that have real, working implementations: rootless Podman's isolation of privileged operations into helper tools [57] and Charliecloud's fully unprivileged model [42].² We close with an analysis of how minimizing container build privileges can be further improved, as well as implications this new build capability has for future HPC workflows.

¹Containers for HPC system services is out of scope for this work.

²Disclosure: Authors R.P. and T.R. are on the Charliecloud team.

2 BACKGROUND

There are two common workflows in use today for container-based HPC application development. The first is for developers to build images on their local laptop or workstation, test on the same hardware, and then transfer the image to an HPC system for production runs. This model allows access to their preferred, local working environment, including editors, tools, configuration, etc. The second is a continuous integration / continuous deployment (CI/CD) approach, where images are automatically built on standalone and isolated resources, such as ephemeral virtual machines, upon code being pushed to a repository and passing tests. In both build workflows, privileged build is a reasonable choice.

However, these workflows raise problems for HPC. First, HPC systems have well-specified architectures that are increasingly diverse, whether a specific x86-64 microarchitecture (e.g., Haswell), a non-x86 CPU family (e.g., ARM or PowerPC), and/or accelerators of some kind (e.g., GPUs). HPC applications are usually performance-sensitive and therefore compiled for the specific architecture of the target supercomputer. On the other hand, developer workstations and CI/CD clouds are not well-specified and must be treated as generic x86-64 resources. This makes it difficult to build HPC applications properly on such resources. While in principle one could cross-compile, in practice this isn't adequately supported by HPC application build systems. Further, testing HPC applications is usually informative only on the target system.

A second problem is resources available only on specific networks or systems. Developers often need licenses for compilers, libraries, or other proprietary code with this limitation. Security-sensitive applications and data often have stringent restrictions on where they may be stored. This means container images must be built and stored in the same places.

In these and other scenarios, building on HPC resources closely matching the target architecture, perhaps supercomputers themselves, properly isolated and with direct access to services such as license servers, is the most natural approach. It can also accelerate development by eliminating the need to transfer images between networks. However, providing container image build capabilities on HPC resources is challenging at present.

The rest of this section explains why. Doing so first requires a deep dive into how containers work, specifically user namespaces and the correspondence between host user IDs (UIDs) or group IDs (GIDs) and those in the container. Based on this knowledge, we introduce a new taxonomy of privilege models for containers. Finally, we detail why naïve unprivileged container build doesn't work, to motivate the more complex approaches later in the paper.

2.1 User namespaces and their implications

Containers are currently implemented using Linux kernel features,³ most importantly *namespaces* [27], which let a process have its own view of certain kernel resources. Crucially, the *mount namespace* gives a process its own mounts and filesystem tree [33], allowing the container to run a different distribution than the host, and the *user namespace* gives a process its own UID and GID space [28, 34].

We give a focused discussion of namespaces, omitting three details: (1) There are about a half dozen other types of namespace. (2) If namespaces are compiled into the kernel, *all* processes are within a namespace of every configured type whether or not containerized, even if some types are disabled via `sysctl`, which governs creation of *new* namespaces. (3) Some namespaces, including `user`, can be nested many levels deep. Our discussion uses a simple two-level host/container division, and we consider only the `user` and `mount` namespaces.

Though incomplete, this focused view is accurate and yields a more accessible discussion. Also, we mostly cover `user` namespaces because they are the key to low-privilege container implementations and have quite subtle behavior.

2.1.1 UID and GID maps. The defining feature of user namespaces is that a process can have UIDs and GIDs inside a namespace that are different from its IDs outside. To accomplish this, the namespace is created with two one-to-one mappings: one between host UIDs and namespace UIDs, and another between host GIDs and namespace GIDs.⁴ Because the map is one-to-one, there is no squashing of multiple IDs to a single ID in either direction. In both cases, it is the host IDs that are used for access control; the namespace IDs are just aliases. For example, one could map the unprivileged invoking host `user` to namespace UID 0 (i.e., `root`).⁵ Benefits of this include (1) giving a user `root` inside a namespace while remaining unprivileged on the host and (2) allocating a set of IDs with well-known numbers (e.g., a distribution's system users and groups) that do not clash with the host. This containerized process would appear to be privileged within the namespace, but in reality is just another unprivileged process. These mappings need not be complete. The four possibilities for a given pair of IDs are:

- (1) *In use on the host, and mapped to a namespace ID.* The namespace ID is simply an alias of the host ID for use within the namespace.
- (2) *Not in use on the host, and mapped to a namespace ID.* This is the same as Case 1, except the host ID happens to not yet be used. The kernel has no notion of "active" UIDs or GIDs, i.e., there is no initialization step for IDs. In particular, files can be owned by these IDs, but outside the namespace they will have no corresponding user or group names.
- (3) *In use on the host, but not mapped to a namespace ID.* These are valid inside the namespace, but there is no way to refer to them. Further, a user with access to a file via an unmapped group can still access it inside the namespace, even though `ls` will show it as `nogroup`. Confusingly, a file owned by a different unmapped group the user *doesn't* have access to will *also* be listed as `nogroup`. System calls and `setuid/setgid` executables do not accept unmapped IDs.
- (4) *Not in use on the host, and not mapped to a namespace ID.* These IDs are not available inside the namespace. Because they are not in use on the host, processes cannot have them

⁴The kernel is concerned only with IDs, which are integers in the range 0 to $2^{32} - 1$ inclusive. Translation to usernames and group names is a user-space operation and may differ between host and container even for the same ID.

⁵In-namespace UID 0 is special because the `execve(2)` system call that transfers control from container runtime to containerized application typically gives the process all capabilities within the namespace [31, 34]. Thus, "UID 0" and "having all capabilities" can be treated the same for our discussion.

³Container implementations on other operating systems work by adding Linux to that operating system, e.g. Docker on Mac's hypervisor approach [44].

Figure 1: Typical privileged UID map for container run by Alice. The file `/etc/subuid` configures the user-space helper for host UIDs Alice and Bob may use; `/proc/self/uidmap` is the subsequent kernel mapping.

when entering the namespace, and they can't be changed to within the namespace.

While correct file ownership is critical in a multi-user system to protect users and processes from one another, multiple users and groups *within an image* is rarely needed for HPC application containers. That is, if an HPC application is containerized, then anyone using the container should be able to see any and all files within it, because that is what's required to run the application. The concept of privileged and unprivileged users is not needed here. Making files accessible for only specific users or groups isn't useful in this context, and could lead to inconsistent behavior depending on who is running the containerized application. This is also true for many containerized HPC services such as databases, which can run as the same user as the jobs they support.

To be clear, there are scenarios where containerized applications or application-support tools do need to run with multiple users or with different privilege levels, e.g. some web services or databases that need to act on behalf of multiple users. However, these are the exception rather than the norm for HPC, and require close sysadmin involvement whether containerized or not.

2.1.2 Privileged ID maps. Processes that hold `CAP_SETUID` or `CAP_SETGID` [31] (and UID 0 typically holds all capabilities) can set up mostly-arbitrary UID and GID maps respectively. This privileged functionality⁶ can be used in otherwise-unprivileged implementations via `setuid` or `setcap` helper programs to set up the user namespace. Currently, the standard helpers are `newuidmap(1)` and `newgidmap(1)` from the `shadow-utils` package [49, 50], though some implementations do provide their own (e.g., Singularity).

A typical configuration for image building is to map the calling user to root, and then map a sequence of unused host UIDs to container UIDs 1– n ; Figure 1 shows an example. GIDs have an

⁶While `CAP_SETUID` is limited privilege, such a process can gain additional privileges by changing its UID to zero, so it's critical to ensure no such execution paths exist.

analogous map. This yields valid in-namespace IDs from 0 to n , with n selected to cover the distribution's system users and groups. Linux's 32-bit IDs help provide enough IDs to everyone who needs to build container images. However, an important task for sysadmins is to ensure that all host IDs really are unused; for example, it's easy to miss IDs on unmounted NFS shares.

The privileged helper approach does need to be implemented carefully, for two reasons. First, the helper is a security boundary, responsible for ensuring that unprivileged users can set up only safe maps; we are aware of real-world errors in such checks, making it a practical as well as theoretical concern. Thus, standard helpers provided by the host OS distribution are preferred, as more rigorous security review and patching is available. Second, sysadmins must correctly configure permissible ID mappings. For example, in Figure 1, if host UID 1001 mapped to container UID 65,537, Alice would have access to all of Bob's files. She could use her container root to `setuid(65537)`, which maps back to host UID 1001, and there are many ways to access the host's filesystem tree from a container. We emphasize the importance of correct configuration here. System management is already complicated, and when new, unfamiliar features are added, the risk of human error increases.

An important corollary of the ID map is that IDs will be correct only within the container. Even though image filesystem trees can be easily available both inside and outside the container, file IDs on the outside will be the mostly-arbitrary host side of the map. For example, images are often stored in tar archives. With privileged ID maps, for correct IDs these archives must be created within the container or use an ID source other than the filesystem.

2.1.3 Unprivileged ID maps. Unprivileged processes can make only limited mappings: (1) host UID to arbitrary in-namespace UID and (2) host GID to arbitrary in-namespace GID [34]. Supplementary groups must remain unmapped. Thus, the process has precisely the same access within the container as on the host.

Because they are unmapped, supplementary groups have reduced functionality: `chgrp(1)` can't be used on them, and they can cause confusion when displayed as `nogroup`.

In other words, for unprivileged user namespaces, the ID map does not matter, except for programs that check for specific IDs. Otherwise, mapped IDs in the unprivileged case are simply a convenience provided by the kernel, with only cosmetic effects.

2.1.4 The `setgroups(2)` trap. This system call requires careful handling when setting up a privileged user namespace due to obscure interactions between the GID map and UNIX permissions. It lets a process add or remove supplementary groups [30], which have different implications in a user namespace.

Adding groups to a process is normally uninteresting from an access control perspective, because it is a privileged system call, i.e. the process already has access to everything. In privileged user namespaces, this operation is more consequential because they can choose whether to enable `setgroups(2)`. If enabled, anyone with root in the namespace also has access to everything protected by all mapped groups. The lesson is the same as the prior section: sysadmins must configure allowed group mappings carefully.

Removing groups from a process has an interesting subtlety. UNIX permissions are evaluated in the order user, group, other — and the first match governs [34]. For example, consider a production

system containing a file `/bin/reboot` owned by `root:managers` and permissions `rwX---r-x`. With this setup, all users *except managers* can reboot the system. However, any manager process able to call `setgroups(2)` can drop managers from its groups, changing the match from group to other, and now that manager can reboot the system. Without user namespaces, this situation again cannot arise, because well-run sites do not give managers root. Privileged helpers are responsible for disabling `setgroups(2)` when acting on behalf of unprivileged users; `newgidmap(1)` did have a vulnerability when it failed to do so [1].

Neither of these issues are concerning for unprivileged user namespaces because `setgroups(2)` is not available [34].

2.2 Container privilege levels

This paper uses the term *privileged* for processes that run with escalated privilege, whether root or non-trivial capabilities [31]. The terms *unprivileged* or *fully unprivileged* refer to processes that run as normal users with no capabilities and *mostly unprivileged* for such processes that call privileged (`setuid` or `setcap`) helper programs; *low privilege* refers to either fully or mostly unprivileged.

We propose a three-level taxonomy of container privilege, based on the presence or absence of the user namespace and whether it is created with host privileges:

- Type I. Mount namespace (or `chroot(2)`) but no user namespace. Privileged setup; shares IDs with the host system. Root inside a container is root on the host.
- Type II. Mount namespace and privileged user namespace. Privileged setup; arbitrarily many UIDs and GIDs independent from the host. Root inside the container is typically mapped to an unprivileged host user, not root.
- Type III. Mount namespace and unprivileged user namespace. Unprivileged setup; only one UID and one GID mapped into the container, which are aliases of the invoking user's IDs. That is, containerized processes remain unprivileged, holding the user's normal unprivileged IDs.

A well-known example of a Type I implementation is Docker [6]. To demonstrate Type II and III, this paper uses rootless Podman [23, 57] and Charliecloud [42] respectively. We expect lessons derived will be applicable to other implementations as well.

2.3 Unprivileged container build is hard

Building a container image largely entails running a sequence of containerized commands on a read-write image. This currently works much like building a system: unpack a Linux distribution base, install the necessary distribution packages (e.g., with `dpkg` or `rpm`), and then install the desired application(s). The difference from *running* containers is that distribution package managers assume privileged access, and key packages need multiple UIDs/GIDs and privileged system calls like `chown(2)` to install. We have not yet encountered any fundamental reason privileged access is needed; the tools install software, rather than using privileged operations like listening on low ports. However, multiple decades of Linux software delivery has done it this way, and it works well. There has not yet been a need to change.

Ideally, we would be able to install packages using a fully unprivileged (Type III) container implementation with no special

```
1| $ cat centos7.dockerfile
2| FROM centos:7
3| RUN echo hello
4| RUN yum install -y openssh
5| $ ch-image build -t foo -f centos7.dockerfile .
6| 1 FROM centos:7
7| 2 RUN ['/bin/sh', '-c', 'echo hello']
8| hello
9| 3 RUN ['/bin/sh', '-c', 'yum install -y openssh']
10| [...]
11| Installing : openssh-7.4p1-21.el7.x86_64
12| Error unpacking rpm package openssh-7.4p1-21.el7.x86_64
13| error: unpacking of archive failed on file [...]: cpio: chown
14| [...]
15| error: build failed: RUN command exited with 1
```

Figure 2: Simple CentOS 7-based Dockerfile that fails to build in a basic Type III container because `chown(2)` failed. We selected the OpenSSH client because it's problematic across distributions and common in HPC user containers. (This and later transcripts have been simplified for clarity. The notation [...] is sometimes used to avoid confusion, but in most cases irrelevant text is simply omitted.)

```
1| $ cat debian10.dockerfile
2| FROM debian:buster
3| RUN echo hello
4| RUN apt-get update
5| RUN apt-get install -y openssh-client
6| $ ch-image build -t foo -f debian10.dockerfile .
7| 1 FROM debian:buster
8| 2 RUN ['/bin/sh', '-c', 'echo hello']
9| hello
10| 3 RUN ['/bin/sh', '-c', 'apt-get update']
11| E: setgroups 65534 failed - setgroups (1: Operation not permitted)
12| E: setegid 65534 failed - setegid (22: Invalid argument)
13| E: seteuid 100 failed - seteuid (22: Invalid argument)
14| [...]
15| error: build failed: RUN command exited with 100
```

Figure 3: Simple Debian 10-based Dockerfile that fails to build in a basic Type III container. Here, `apt-get` encountered errors trying (ironically) to drop privileges.

modifications or alterations. Unfortunately, this is not the case. For example, Figure 2 presents a simple CentOS-based Dockerfile that fails to build in such a configuration. In this case, the system call `chown(2)` fails because while the process appears to be root within the container, it is really an unprivileged process running as the invoking user, and thus may not call `chown(2)` [32].

This problem is not unique to RPM-based distributions. In Figure 3, a Debian-based Dockerfile also fails to build. In this case, `apt-get` tries to drop privileges and change to user `_apt` (UID 100) to sandbox downloading and external dependency solving [14, 59]. This yields a series of failed system calls. `setgroups(2)` fails because that system call is not permitted in an unprivileged container [34]. `setresgid(2)` (not `setegid(2)` and `seteuid(2)` as in the error messages) fails twice because again `apt-get` is actually unprivileged and may not make those system calls [29].

The colliding assumptions of (1) package managers require root and (2) unprivileged containers disallow root mean that unprivileged container build is a harder problem than unprivileged run. The next section outlines solutions.

3 APPROACHES FOR CONTAINER BUILD

Above, we’ve motivated the need for unprivileged container build in HPC and detailed how user namespaces can provide reduced-privileged containers. In this section, we take a step back, assessing existing container implementations used in HPC and approaches for limiting build privilege. This assessment reinforces our conclusion that user namespace-based container builds are currently the lowest-cost path forward for HPC.

3.1 Container implementations used in HPC

There are several container implementations in active or potential use for HPC. Containers originated in the web applications world [39], which often uses virtual machine sandboxes as noted above, so running applications in a Type I container is a reasonable choice whether or not the containerized processes are privileged. Also, user namespaces were not available until Linux 3.8, released in February 2013 [55], and they were not fully supported in Red Hat Enterprise Linux (RHEL) or its derivatives, upon which most HPC centers rely, until version 7.6 in October 2018 [45]. Without user namespaces, only Type I containers are possible.

Docker was the first container implementation to become widely popular. Its initial public release was March 2013 and supported Linux 2.6.24 [6], making it Type I by necessity. Even simply having access to the “docker” command is equivalent to root “by design” [46]. Another key design decision was its client-daemon execution model [23]. Processes started with “docker run” are descendants of the Docker daemon, not the shell, which is undesirable for HPC because it is another service to manage/monitor, breaks process tracking by resource managers, and can introduce performance jitter [17]. Docker quickly became the industry standard, and in fact many containers intended for supercomputers are currently built with Docker using the workflows noted above. Because its approach worked fine for web applications, it gained inertia and mind share even as unprivileged kernel container features matured. Docker did add a Type II mode in 2019 [15], though it is not the default and to our knowledge is not yet widely used in practice.

The HPC community has also produced container implementations, some with build capability. The most popular HPC container implementation is currently Singularity [35], which can run as either Type I or II (branded “fakeroot”). As of this writing, Singularity 3.7 can build in Type II mode, but only from Singularity definition files. Building from standard Dockerfiles requires a separate builder (e.g., Docker) followed by conversion to Singularity’s image format, which is a limiting factor for interoperability. Type I HPC examples include both Shifter [24] and Sarus [7], though currently these focus on distributed container launch rather than build. Another implementation of interest is Enroot, which advertises itself as “fully unprivileged” [3] with “no setuid binary” [9], i.e., Type III. However, as of the current version 3.3, it does not have a build capability, relying on conversion of existing images.

3.2 Summary of build options

In this section, we propose some requirements for container build and analyze three general approaches in light of them. First, the build recipe (typically a Dockerfile) should require no modifications. Because existing recipes are often used in whole or in part,

modifications impact productivity and portability, which are major benefits of containers. Second, we treat the current behavior of distribution packaging tools as unchangeable. (We discuss below in §6.2.3 how these tools’ privilege needs could be relaxed.) With these requirements in mind, we consider three approaches:

- (1) *Sandboxed build system.* One way to work around increased privileges is to create an isolated environment specifically for image builds. Many approaches are available, including most commonly virtual machines or bare-metal systems with no shared resources such as production filesystems. When creating these environments, the risk of users circumventing controls must be carefully minimized. These models do work well with CI and other automated approaches, where further limits may be available.
- (2) *Type II containers.* With the isolated UIDs and GIDs provided by user namespaces, this approach can provide a mostly-seamless illusion of root within the container while remaining mostly unprivileged. Privileged helper tools are needed to set up the user namespace, and sites must maintain correct mapping files to ensure users are properly isolated from each other.
- (3) *Type III containers.* This type of container is fully unprivileged, but limited UIDs and GIDs mean more is needed to install arbitrary packages. One solution is a wrapper like fakeroot (1), which intercepts privileged system calls and fakes their success. This wrapper can be injected automatically to meet the first requirement above.

Recall that building containers with low privilege on HPC resources is the focus of this work. Option 1 offers a pragmatic way to work around privilege issues and is already in use at many sites, e.g., via the Sylabs Enterprise Remote Builder [53]. However, isolated build environment may not be able to access needed resources, such as private code or licenses. Thus, our view is that Options 2 and 3 are the most promising.

These two options are distinguished by user namespaces being privileged or unprivileged, respectively. Aside from details of any specific implementation, there are two key factors to assess this distinction. First, are the security challenges of privileged namespace setup helpers and their configuration acceptable to a site? Second, how much does a site care that container images include specific or multiple IDs? If the helpers are acceptable and specific/multiple IDs are needed for exact container images, then sites should consider Option 2, detailed below in §4 using rootless Podman. Otherwise, they should consider Option 3, detailed in §5 using Charliecloud.

4 PODMAN

As detailed above, one way to provide a reasonable container build capability for HPC resources is by using a Type II container build solution directly on a supercomputer. A recent container implementation under this category is Podman,⁷ and in particular the “rootless Podman” configuration, which is particularly relevant to HPC [56, 57]. The main design goals of rootless Podman are to have the same command-line interface (CLI) as Docker, remove root-level privileges, and use a fork-exec model rather than Docker’s client-daemon model [23]. Because Podman is CLI equivalent to Docker,

⁷<https://podman.io>

many users can successfully utilize Podman by alias `docker=podman` and use as expected. Podman also adheres to the OCI spec for container compatibility and interoperability, making transition to Podman from existing container implementations effectively seamless. While Podman is open source, it is heavily supported by Red Hat and integrated into their Linux distribution as of RHEL7.8.

While we refer to rootless Podman, or Podman more generally as a Type II container implementation in this paper, the implementation is in fact based on the same code as the Buildah tool and other related container utilities. Podman in this sense only provides a CLI interface identical to Docker, whereas Buildah provides more advanced and custom container build features. As Podman and Buildah leverage the same codebase for build operations, we can assume they are functionally equivalent in this paper for purposes of brevity. As Podman avoids Docker's daemon design, it has gained significant interest for HPC applications as it is a much better fit for HPC deployments where daemons are generally discouraged whenever possible. While this design decision can seem inconsequential, the removal of a daemon can have significant improvements to overall system performance and manageability [41].

4.1 Rootless Podman with privileged helpers

As detailed in Section 2.1, the ability to minimize privilege with rootless Podman is a key feature for HPC. In actuality, rootless Podman refers to the ability to build and run containers without administrative privileges. While there are other potential configurations, the most common implementation currently uses Privileged ID maps to help mitigate the security issues of Docker's inherent privilege escalation. Effectively, this is a Type II configuration. Podman itself remains completely unprivileged; instead a set of carefully managed tools provided by shadow-utils are executed by Podman to set up user namespaces mappings. The shadow-utils executables, `newuidmap` and `setgidmap`, map the specified UIDs and GIDs into the container's user namespace. These tools read `/etc/subuid` and `/etc/subgid`, respectively, which outline the user namespace mappings. The shadow-utils executables are installed using `CAP_SETUID`, which helps minimize risk of privilege escalation compared to using a `SETUID` bit. While this Linux capability still poses a potential security risk, the multiple points of privilege separation between Podman and shadow-utils, along with the careful validation of the executables' implementation by a trusted Linux OS distribution (such as Red Hat), can reasonably minimize the possibility of escalation to the point where we believe Podman can be safely used in a production environment. Furthermore, Podman can be additionally secured with SELinux to further prevent potential points of exploitation.

An example namespace mapping is provided in Figure 4. Here, you can see the user `alice` is able to allocate 65535 UIDs, starting at UID 200000. The user namespace mapping definitions cannot exceed the maximum user namespaces as configured in `/proc/sys/user/max_user_namespaces`. While the approach of privileged helpers works well in practice, it needs to be implemented carefully as the `setuid` helper is responsible for enforcing a security boundary to ensure unprivileged users can set up only safe maps. These mappings need to be specified by the administrator upon Podman installation for existing users. Newer versions of shadow-utils can

```
1| $ cat /etc/subuid
2| # USER : STARTUID : TOTALUIDs
3| alice:200000:65536
4| bob:300000:65536
5| $ podman unshare cat /proc/self/uid_map
6|      0      1234      1
7|      1    200000    65536
```

Figure 4: An example subuid file outlining the user namespace mappings used by Podman, followed by a listing of the UID map used by Podman. The unshare command illustrates how user namespace mapping is mapping UID 0 in the container to 1234 on the host for a range of 1 UIDs. While not listed, subgid file would have similar mappings.

automatically manage the setup using `useradd` and `usermod --add-subuids`, which further simplifies administrative burden and lowers the risk of conflicts. This ability for system administrators to carefully control user access to Podman's capabilities can be an asset for facilities looking to first prototype rootless Podman without fully enabling the utility system-wide.

Podman also provides several additional features to provide a full-featured and production container solution to users. First, Podman by default leverages the runc container runtime for OCI compliance. This usage of runc as the underlying base runtime is similar to the implementation of Sarus [7] and helps conform to OCI standards to provide interoperability and compatibility with not only other container implementations, but also many OCI-compliant container registry services. Second, Podman uses the fuse-overlayfs storage driver which provides unprivileged mount operations using a fuse-backed overlay file-system. Podman can also use the VFS driver, however this implementation is much slower and has significant storage overhead, and should only be used when absolutely necessary with existing installations (eg: RHEL7).

With rootless Podman, cgroups are left unused as cgroup operations by default are generally root-level actions. This is a convenient coincidence for HPC, since the on-node resource management is typically handled by the resource manager and job scheduler and having multiple services trying to interact with cgroups should be avoided. There still may be future utility in cgroups operations, so prototype work is underway to implement cgroups v2 in userspace via the crun runtime, which enables cgroups control in a completely unprivileged context. In summary, rootless Podman with shadow-utils privileged mappers to help manage user namespace mappings, provide a full-featured and Docker CLI-equivalent container build solution on HPC resources. When Podman is configured appropriately, the examples detailed in Figures 2 and 3 will both succeed as expected when executed by a normal, unprivileged user.

4.1.1 Unprivileged Podman. While the default rootless Podman usage is designed around the privileged mappers from shadow-utils to provide a comprehensive user namespace mapping, it is also possible to set up Podman in an experimental unprivileged mode. Instead of the namespace mappers as illustrated in Figure 4, the subuid mappers are discarded and a single UID is mapped to the container. When the UID mappers for a given user are left unset and the `--ignore_chown_errors` option is enabled with either the VFS or overlay driver, Podman then operates with unprivileged


```

1| $ cat /etc/subuid
2| $ podman unshare cat /proc/self/uid_map
3|           0       1234       1

```

Figure 5: Podman UID mapping in unprivileged mode. This setup does not use privileged helpers but has limitations.

Figure 6: Container Build Workflow on Astra, with Podman

namespaces with one UID mapping, as illustrated in Figure 5. However, the `yum install openssh-server` example in Section 2 will fail because `/proc` and `/sys` mappings in the container are owned by user `nobody`; a consequence of the namespace mapping. While further workarounds with `fakeroot(1)` as detailed in the following section could be applied here, this setup is not yet used in practice.

4.2 Podman on the Astra Supercomputer

To enable user-driven container build capabilities directly on HPC resources, rootless Podman was first deployed and tested on the Astra supercomputer [40]. Astra has proven to be an ideal first system to prototype the use of Podman for two reasons. First and most critically, Astra was the first Arm-based supercomputer on the Top 500 list, which meant existing containers based on x86_64 would not execute on Astra. Instead, users had an immediate need to *build* new container images specifically for the aarch64 ISA, and in particular the Marvell Thunder X2 CPUs. This requirement included a containerizing ATSE, the Advanced Tri-Lab Software Environment, which provides an open, modular, extensible, community-engaged, and vendor-adaptable software ecosystem for which many production HPC codes from NNSA labs depend on [62]. Second, Astra was initially deployed as a prototype system, which gave researchers additional latitude to test and evaluate experimental system software capabilities like those found with Podman.

With Astra, Podman was used as part of a multi-stage container DevOps process and is outlined in Figure 6. First, `podman build` is invoked from the Astra login node to construct the ATSE container

build process in a directory which includes the Dockerfile for ATSE. This builds all of the ATSE components, including compilers, MPI libraries, third-party libraries, and test applications all directly in a container on the login node. This command is invoked by the user. Once the build is completed, it can be pushed to a OCI-compliant container registry (in this case the Gitlab Container Registry Service). A container registry is important to leverage in this workflow as it provides persistence to container images which could help in portability, debugging with old versions, or general future reproducibility. Finally, the container image built on the supercomputer can be deployed in parallel using the local resource management tool and an HPC container runtime. This was originally demonstrated with Singularity, however any HPC container runtime such as Charliecloud or Shifter could also be used with a similar step. While Podman could potentially be used for parallel execution [19], it has not yet been optimized for distributed container execution and likely would introduce significant overhead at scale.

To our knowledge, Astra’s usage of Podman version 1.4.4 (Buildah 1.9.0) is the first demonstration of unprivileged users building OCI containers directly on a supercomputer. However, the implementation is not without limitations. First, Podman with RHEL7.6 used the VFS driver, which is known to add overhead. Newer versions of Podman, as found in RHEL8, can leverage the overlayfs-fuse driver for better performance. Second, the UID/GID mappers cannot work when the container storage location is a shared filesystem, such as NFS. The reason is that the filesystem server has no way to enforce the file creation of different UIDs on the server side and cannot utilize the user namespace configuration. For Astra, either `/tmp` or local disk can be used for container storage on the login nodes, which requires additional configuration by users and does not enable any distributed launch abilities.

5 CHARLIECLOUD

Charliecloud⁸ is a container implementation designed for HPC applications with emphasis on the unprivileged use cases discussed in the first half of this paper [42]. It is distinguished from other implementations because it is lightweight (roughly 4,000 lines of code as tested, version 0.23) and has been Type III from its first release. Its runtime `ch-run(1)` is written in C, and it has a Dockerfile interpreter and container registry API client written in Python, `ch-image(1)`. To our knowledge, this is the first fully unprivileged builder able to build Dockerfiles that make privileged system calls, which it does by automatically injecting `fakeroot(1)` into the build. This section details the approach. We first present background on `fakeroot(1)` and show how it can work around the limitations of unprivileged container build; we then describe Charliecloud’s injection of the wrapper, which should be transferrable to any Type III implementation.

5.1 Background: Faking root with fakeroot(1)

`fakeroot(1)` is a program to run a command in an environment that appears to be privileged but is not [12].⁹ It does this by intercepting privileged and privileged-adjacent system calls and lying

⁸<https://hpc.github.io/charliecloud>

⁹Not to be confused with Singularity’s Type II mode, also called “fakeroot”.

```

1| $ fakeroot ./fakeroot.sh
2| + touch test.file
3| + chown nobody test.file
4| + mknod test.dev c 1 1
5| + ls -lh test.dev test.file
6| crw-r----- 1 root  root  1, 1 Feb 10 18:09 test.dev
7| -rw-r----- 1 nobody root   0 Feb 10 18:09 test.file
8| $ ls -lh test*
9| -rw-r----- 1 alice  alice  0 Feb 10 18:09 test.dev
10| -rw-r----- 1 alice  alice  0 Feb 10 18:09 test.file

```

Figure 7: Example of fakeroot(1) use. This script changes ownership of a file and then creates a device file, both of which are privileged operations. fakeroot(1) intercepts the system calls and causes them to “succeed”. Within the fakeroot(1) context, ls(1) shows the expected results (device file and nobody-owned file, respectively); the subsequent unwrapped ls(1) exposes the lies.

to the wrapped process about their results. For example, an unprivileged process could chown(2), and fakeroot(1) will return success from its own implementation without ever actually calling the system call. It also remembers which lies it told, to make later intercepted system calls return consistent results. In this example, fake results for stat(2) include the user and/or group set by the earlier fake chown(2). Because stat(2) is not privileged, the wrapper really will make the system call, but then adjust the results to make them look correct in the “privileged” context. Figure 7 shows an example of fakeroot(1) use.

fakeroot(1) is not a perfect simulation but rather just enough to work for its intended purpose, which is building distribution packages, allowing “users to create archives (tar, ar, .deb etc.) with files in them with root permissions/ownership” [12]. The focus is on filesystem metadata; e.g., one can’t use fakeroot(1) to listen on an unprivileged port and make a process see a privileged port. Fortunately, this purpose is highly congruent with the needs of unprivileged container build.

There are three fakeroot(1) implementations we are aware of, summarized in Table 1. They each have different quirks; for example, LD_PRELOAD implementations are architecture-independent but cannot wrap statically linked executables, while ptrace(2) are the reverse. We’ve encountered packages that fakeroot cannot install but fakeroot-ng and pseudo can; fakeroot appears to be under more active development.

5.2 Fully unprivileged build via fakeroot(1)

We can use fakeroot(1) to fix the failed container builds in §2.3. This approach will squash the actual ownership of all files installed to the invoking user, but that usually does not matter for HPC applications, and any downstream Type III users that pull the image will change ownership to themselves anyway, like tar(1) [25].

Figure 8 shows modifications of the CentOS 7 Dockerfile from Figure 2 to build it under ch-image’s simplest configuration. Three changes were needed:

- (1) Install EPEL, because CentOS has no fakeroot package in the default repositories.
- (2) Install fakeroot.
- (3) Prepend fakeroot(1) to the openssh install command.

```

1| $ cat centos7-fr.dockerfile
2| FROM centos:7
3| RUN yum install -y epel-release
4| RUN yum install -y fakeroot
5| RUN echo hello
6| RUN fakeroot yum install -y openssh
7| $ ch-image build -t foo -f centos7-fr.dockerfile .
8| 1 FROM centos:7
9| 2 RUN ['/bin/sh', '-c', 'yum install -y epel-release']
10| [...]
11| Complete!
12| 3 RUN ['/bin/sh', '-c', 'yum install -y fakeroot']
13| [...]
14| Complete!
15| 4 RUN ['/bin/sh', '-c', 'echo hello']
16| hello
17| 5 RUN ['/bin/sh', '-c', 'fakeroot yum install -y openssh']
18| [...]
19| Complete!
20| grown in 5 instructions: foo

```

Figure 8: CentOS 7 Dockerfile from Figure 2, modified to build successfully by wrapping the offending yum install with fakeroot(1).

```

1| $ cat debian10-fr.dockerfile
2| FROM debian:buster
3| RUN echo 'APT::Sandbox::User "root";' >
4|   /etc/apt/apt.conf.d/no-sandbox
5| RUN echo hello
6| RUN apt-get update
7| RUN apt-get install -y pseudo
8| RUN fakeroot apt-get install -y openssh-client
9| $ ch-image build -t foo -f debian10-fr.dockerfile .
10| 1 FROM debian:buster
11| 2 RUN ['/bin/sh', '-c', 'echo \'APT::Sandbox::User "root";\' >
12|   /etc/apt/apt.conf.d/no-sandbox']
13| 3 RUN ['/bin/sh', '-c', 'echo hello']
14| hello
15| 4 RUN ['/bin/sh', '-c', 'apt-get update']
16| [...]
17| Fetched 8422 kB in 7s (1214 kB/s)
18| Reading package lists...
19| 5 RUN ['/bin/sh', '-c', 'apt-get install -y pseudo']
20| [...]
21| Setting up pseudo (1.9.0+git20180920-1) ...
22| [...]
23| W: chown to root:adm of file /var/log/apt/term.log failed [...]
24| 6 RUN ['/bin/sh', '-c', 'fakeroot apt-get install -y
25|   openssh-client']
26| [...]
27| Setting up openssh-client (1:7.9p1-10+deb10u2) ...
28| Setting up libxext6:amd64 (2:1.3.3-1+b2) ...
29| Setting up xauth (1:1.0.10-1) ...
30| Processing triggers for libc-bin (2.28-10) ...
31| grown in 6 instructions: foo

```

Figure 9: Debian 10 Dockerfile from Figure 3, modified by configuring apt-get(1) to not drop privileges and wrapping the offending apt-get(1) calls with fakeroot(1).

The first two install steps do use yum(1), but fortunately these invocations work without fakeroot(1).

Figure 9 shows analogous modifications of the Debian 10 build from Figure 3. This needed more complex changes:

- (1) Disable apt-get(1)’s privilege dropping.
- (2) Update the package indexes. The base image contains none, so no packages can be installed without this update.
- (3) Install pseudo. (In our experience, the fakeroot package in Debian 10 was not able to install the packages we tested.)

implementation	initial release	latest version	approach	architectures	daemon?	persistence
fakeroot [11, 12, 13]	1997-Jun	2020-Oct (1.25.3)	LD_PRELOAD	any	yes	save/restore from file
fakeroot-ng [48, 51]	2008-Jan	2013-Apr (0.18)	ptrace(2)	PPC, x86, x86-64	yes	save/restore from file
pseudo [2, 47]	2010-Mar	2018-Jan (1.9.0)	LD_PRELOAD	any	yes	database

Table 1: Summary of fakeroot(1) implementations.

```

1| $ ch-image build --force -t foo -f centos7.dockerfile
2|   1 FROM centos:7
3| will use --force: rhel7: CentOS/RHEL 7
4|   2 RUN ['/bin/sh', '-c', 'echo hello']
5| hello
6|   3 RUN ['/bin/sh', '-c', 'yum install -y openssh']
7| workarounds: init step 1: checking: $ command -v fakeroot >
   | ↪ /dev/null
8| workarounds: init step 1: $ set -ex; if ! grep -Eq '\[epel\]'
   | ↪ /etc/yum.conf /etc/yum.repos.d/*; then yum install -y
   | ↪ epel-release; yum-config-manager --disable epel; fi; yum
   | ↪ --enablerepo=epel install -y fakeroot;
9| + grep -Eq '\[epel\]' [...]
10| + yum install -y epel-release
11| [...]
12| Complete!
13| + yum-config-manager --disable epel
14| [...]
15| + yum --enablerepo=epel install -y fakeroot
16| [...]
17| Complete!
18| workarounds: RUN: new command: ['fakeroot', '/bin/sh', '-c', 'yum
   | ↪ install -y openssh']
19| [...]
20| Complete!
21| --force: init OK & modified 1 RUN instructions
22| grown in 3 instructions: foo

```

Figure 10: Successful CentOS 7 build with unmodified Dockerfile. The fakeroot(1) calls are almost identical to Figure 8 but automatically injected by ch-image.

(4) Prepend fakeroot(1) to the install of openssh-client.

Line 21 still has a warning. In our experience, Debian package management under fakeroot(1) occasionally complains about privileged operations failing, but these warnings do not stop the build.

These modifications leave us with working builds, but we want to build *without* modifying any Dockerfiles.

5.3 Auto-injection of fakeroot(1)

Unprivileged build of unmodified Dockerfiles can be done by automatically modifying the build. Charliecloud’s ch-image(1) contains such an implementation. Its design principles are:

- (1) Be clear and explicit about what is happening.
- (2) Minimize changes to the build.
- (3) Modify the build only if the user requests it, but otherwise say what *could* be modified.

We outline Charliecloud’s implementation in the context of the two running example Dockerfiles.

5.3.1 Example: CentOS 7. We can build the CentOS 7 Dockerfile in Figure 2 with “ch-image --force” to enable the modifications; Figure 10 shows the result. Like the other figures, this transcript shows the in-container commands that actually execute the RUN

instructions, as well as new commands to implement the modifications. The basic steps are simple: (1) test if the distribution is supported, (2) install fakeroot(1) if needed, and (3) augment RUN instructions that seem to need it. This Dockerfile works as follows.

First, the FROM instruction (lines 2–3) tells us that ch-image can use modifications requested with --force, though it doesn’t yet know whether it *will*. It tested known configurations and found a matching one called rhel7: the file /etc/redhat-release exists and its contents match the regular expression “release 7\.”. (This approach avoids executing a command within the container.)

Next, ch-image executes the first RUN instruction (lines 4–5) normally, because it doesn’t seem to need modification. The second RUN instruction (line 6), contains the string “yum”, a keyword configured to trigger modification of the instruction.

Because this is the first modified RUN, ch-image must initialize fakeroot(1). A configuration has a sequence of steps to do this (though for rhel7 there is only one step), and each step contains two phases: a shell command to check if the step needs to be done, and a shell command to do the step. Here, the check (line 7) tests if fakeroot(1) is already installed in the image. In this example, fakeroot(1) is *not* already present, so ch-image proceeds to the installation step (line 8), which for rhel7 is a long pipeline with individual commands echoed for clarity (lines 9, 10, 13, 15). The pipeline installs EPEL if it’s not already installed, but doesn’t enable it because EPEL can cause unexpected upgrades of standard packages, and then it installs fakeroot from EPEL. Files are grep(1)ed directly, rather than using yum repolist, because the latter has side effects, e.g. refreshing caches from the internet.

Finally, the instruction is modified by inserting “fakeroot” before the /bin/sh call that executes it. This would be repeated for later modifiable RUNs, if there were any.

If the user had *not* specified --force, some of this would still occur. ch-image still looks for a matching configuration, and the keyword test is still done for RUN instructions. The result is that ch-image knows that that modifications are *available*, which it suggests if the build fails. In fact, it suggested --force in a transcript line omitted from Figure 2.

5.3.2 Example: Debian 10. Figure 11 shows a successful build of the Dockerfile in Figure 3 with ch-image --force;. In this case, ch-image selected configuration debderiv (line 3) because the file /etc/os-release contains the string “buster”. The first RUN is again left unmodified (lines 4–5).

The second RUN is deemed modifiable because it contains the string “apt-get”. Initialization for Buster takes two steps. First, ch-image tests whether the APT sandbox is disabled by configuration or if user _apt is missing (line 7); it’s not, so a configuration

```

1| $ ch-image build --force -t foo -f debian10.dockerfile
2| 1 FROM debian:buster
3| will use --force: debderiv: Debian (9, 10) or Ubuntu (16, 18, 20)
4| 2 RUN ['/bin/sh', '-c', 'echo hello']
5| hello
6| 3 RUN ['/bin/sh', '-c', 'apt-get update']
7| workarounds: init step 1: checking: $ apt-config dump | fgrep -q
  ↳ 'APT::Sandbox::User "root"' || ! fgrep -q _apt /etc/passwd
8| workarounds: init step 1: $ echo 'APT::Sandbox::User "root";' >
  ↳ /etc/apt/apt.conf.d/no-sandbox
9| workarounds: init step 2: checking: $ command -v fakeroot >
  ↳ /dev/null
10| workarounds: init step 2: $ apt-get update && apt-get install -y
  ↳ pseudo
11| [...]
12| Setting up pseudo (1.9.0+git20180920-1) ...
13| [...]
14| workarounds: RUN: new command: ['fakeroot', '/bin/sh', '-c',
  ↳ 'apt-get update']
15| [...]
16| 4 RUN ['/bin/sh', '-c', 'apt-get install -y openssh-client']
17| workarounds: RUN: new command: ['fakeroot', '/bin/sh', '-c',
  ↳ 'apt-get install -y openssh-client']
18| [...]
19| --force: init OK & modified 2 RUN instructions
20| grown in 4 instructions: foo

```

Figure 11: Successful Debian 10 build with unmodified Dockerfile. The `fakeroot(1)` calls are almost identical to Figure 9 but automatically injected by `ch-image`.

file (line 8) to disable it is added. Next, it tests for the availability of `fakeroot(1)` (line 9); it's not present, so check whether the package `pseudo` (line 10) is installed, which is in the standard repositories. It would be better not to update the package indexes with `apt-get update` because this is an operation with many side effects, but the base images ship with no indexes, so nothing can be installed until after this command.

The modified RUN is then executed (line 14); note that `ch-image` is not smart enough to notice that it's now redundant and could have been skipped. Finally, modify and execute the RUN which installs the problem package `openssh-client` (line 16–17).

5.3.3 Production. `ch-build --force` also works for various production applications at Los Alamos. For example, one application has integrated Charliecloud container build into its CI pipeline using a sequence of three Dockerfiles: the first installs and configures OpenMPI in a CentOS base image, the second installs the complex Spack [18] environment needed by the application, and the third builds the application itself. Once built, the third image is pushed to a private container registry. The validation stage then pulls this image and uses it to run the application's test suite. Build and validate both run on supercomputer compute nodes using normal jobs, and the pipeline is coordinated by a separate GitLab server.

6 DISCUSSION

We have argued that HPC should minimize privilege for container image build, and we have detailed how container implementations can reduce privilege with user namespaces. We demonstrated two approaches with real products: mostly-unprivileged Type II rootless Podman and fully-unprivileged Type III Charliecloud. We now compare the two approaches, outline future work, and discuss how these methods for container builds can impact the field of HPC more broadly.

6.1 Mostly vs. fully unprivileged build

To demonstrate Type II container build via privileged user namespaces, we used rootless Podman. The main advantage of this approach is that it retains file ownership in the container image and presents no UID/GID oddities to distribution tools, allowing traditional builds to work for otherwise-unprivileged users. The main disadvantage is reliance on correct implementation and configuration of the privileged helper tools. One current problem with rootless Podman is that its user namespaces implementation clashes with shared filesystems, including default-configured Lustre, GPFS, and NFS. The issue is not a container problem per se, but rather the user extended attributes (xattrs) Podman uses for its ID mappings.

To demonstrate Type III container build via unprivileged user namespaces, we used Charliecloud. The main advantage here is that the entire build process is fully unprivileged; all security boundaries remain within the Linux kernel. However, the Charliecloud build process is more complex; current disadvantages include:

- (1) The `fakeroot(1)` wrapper is required for a large fraction of builds, which introduces another layer of indirection. While this can be automated, the wrapper must be applied separately to each RUN instruction, and it's extremely difficult to tell precisely which RUNs really need it.
- (2) The resulting image is slightly different. All files must be owned by a single user (the unprivileged user who did the build), and the image cannot contain privileged special files such as devices. On push, Charliecloud changes ownership for all image files to `root:root` and clears `setuid/setgid` bits, to avoid leaking site IDs. Charliecloud currently adds two further complications: `fakeroot(1)` is installed into the image, and images are single-layer, in contrast to other implementations that push images as multiple layers. In general, however, these differences should not impact the behavior of HPC applications.
- (3) Charliecloud lacks a per-instruction build cache, in contrast to other leading Dockerfile interpreters including Podman and Docker. This caching can greatly accelerate repetitive builds, such as during iterative development.

On balance, neither solution currently fits all HPC use cases. Generally, if preserving file ownership within container images is more important than the security cost of privileged helper tools, Type II builders like Podman are likely a good choice; if the opposite, then Type III builders like Charliecloud are likely preferred.

6.2 Future work

Both Type II and Type III container build have high-quality, actively developed implementations available today. However, there are still several opportunities for improvement.

6.2.1 Recommendations for Type II implementations. First, adding a Type III mode to Type II implementations that already use user namespaces is straightforward; doing so would make an implementation more flexible and adaptable. Second, shared filesystems are critical for HPC, so the xattrs support is still necessary. For Lustre, support needs to be enabled on both the metadata server and storage targets. For NFS, changes in Linux 5.9 along with NFSv4 [38] server bring in xattrs support, though at the time of writing these

changes have yet to be extended into distributions common in HPC. We have not evaluated GPFS support at the time of writing. To our knowledge, none of these configurations have yet been tested with Type II containers.

6.2.2 Recommendations for Type III implementations. In addition to Charliecloud-specific improvements in progress like layers and build caching, we identify three more general refinements here:

- (1) Fix fakeroot(1). Not all implementations can install all packages; characterize the scope of the problem and address it. With a sufficiently robust fakeroot(1), Type III could become a drop-in replacement for Type I, eliminating the need for the Type II privilege compromise.
- (2) Preserve file ownership. fakeroot(1) does track this so it can tell consistent lies; the information could be extracted and used when exporting or pushing images. For example, Charliecloud uses the Python tarfile package and could create layer archives that reflect fakeroot(1)'s database rather than the filesystem.
- (3) Move fakeroot(1). Rather than installing in the image itself, the wrapper could be moved into the container implementation. This would simplify it and also ease the prior item. At least one fakeroot(1) implementation already has a libfakeroot [12].

6.2.3 Recommendations for distributions. Beyond improvements to container build implementations, package managers and other distribution tools could improve support for unprivileged users, eliminating the need for Type II privileged code or Type III wrappers. Given increasing use of containers across the computing industry, there may be compelling reasons to support them even beyond HPC. The complexity of this effort should be better understood, and edge cases and older distributions mean container build tools will need their own low-privilege build support for some time.

6.2.4 Recommendations for ID maps. Consider that fakeroot(1) manages user identity and file ownership in user space; in contrast, the user namespace ID maps rely on the kernel to track this state. Currently, Type II kernel ID maps require helper tools like newuidmap(1) and newgidmap(1); they are privileged because this is a security-sensitive operation and to ensure system configuration (/etc/subuid and /etc/subgid) is respected. Future kernel versions could provide other mechanisms to expand the utility of unprivileged maps. For example, supplemental groups could become mappable, general policies could be implemented such as “host UID maps to container root and guaranteed-unique host UIDs map to all other container UIDs”, or the kernel could manage the fake ID database with actual files stored as the invoking user.

6.2.5 Recommendations for standards. Above, we argued that file ownership for HPC application containers is usually not needed, and instead it's mostly an artifact of legacy distribution tools used for building images. A flattened file tree where all users have equivalent access, like that produced by Charliecloud or Singularity's SIF image format, is sufficient and in fact advantageous for most HPC applications. We suspect many non-HPC container use cases have

similar requirements. A potential extension to the OCI specification and/or the Dockerfile language is explicit marking of images to disallow, allow, or require them to be ownership-flattened.

6.3 Impact on HPC

This paper argues that low-privilege container build will have significant impact for the science enabled by HPC, because these capabilities allow containers to be built directly on the HPC resources where they will be run. This property is especially important for HPC developers and users due to the specialized nature of supercomputing hardware, which makes building on dissimilar laptops, workstations, or CI/CD virtual machines undesirable.

We expect these capabilities to also improve DevOps modernization of scientific computing application development and deployment. Specifically, DevOps coupled with low-privilege container build will allow CI/CD pipelines to execute directly on supercomputing resources such as login/front-end and compute nodes, perhaps in parallel across multiple supercomputers or node types to automatically produce specialized container images. For future exascale supercomputers in particular, this user-based build capability, coupled with automated CI runners will enhance development, testing, and deployment of HPC applications directly on the appropriate resources. This integration should reduce not only the effort of developing or porting complex codes on new systems, but also that of managing development resources matching the architecture of one site's supercomputer across many other sites.

ACKNOWLEDGMENTS

Eduardo Arango, Dan Walsh, Scott McCarty, and Valentin Rothberg of Red Hat, Inc. helped us understand rootless Podman. Peter Wienemann of the Debian Project clarified APT privilege dropping. Michael Kerrisk of man7.org Training and Consulting provided corrections for the namespaces discussion. (All remaining errors are of course ours alone.) Kevin Pedretti and the ATSE team helped implement Astra's container workflow.

This work was supported in part by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy (DOE) Office of Science and the National Nuclear Security Administration (NNSA); the Advanced Simulation and Computing Program (ASC); and the LANL Institutional Computing Program, which is supported by the U.S. DOE's NNSA under contract 89233218CNA000001. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. DOE's NNSA under contract DE-NA0003525. LA-UR 21-23314; SAND2021-4332 O.

REFERENCES

- [1] 2018. CVE-2018-7169 Detail. <https://nvd.nist.gov/vuln/detail/CVE-2018-7169>
- [2] 2020. pseudo(1). <https://manpages.debian.org/testing/pseudo/pseudo.1.en.html>
- [3] Felix Abecassis and Jonathan Calmels. 2020. Distributed HPC applications with unprivileged containers. In *Proc. FOSDEM*. https://archive.fosdem.org/2020/schedule/event/containers_hpc_unprivileged/attachments/slides/3711/export/events/attachments/containers_hpc_unprivileged/slides/3711/containers_hpc_unprivileged.pdf
- [4] Subil Abraham, Arnab K. Paul, Redwan Ibne Seraj Khan, and Ali R. Butt. 2020. On the use of containers in high performance computing environments. In *Proc.*

- Cloud Computing*. <https://doi.org/10.1109/CLOUD49709.2020.00048>
- [5] Brian Austin et al. 2020. NERSC-10 workload analysis. Presentation slides. https://portal.nersc.gov/project/m888/nersc10/workload/N10_Workload_Analysis.latest.pdf
 - [6] Abel Avram. 2013. Docker: Automated and consistent software deployments. *InfoQ* (March 2013). <https://www.infoq.com/news/2013/03/Docker/>
 - [7] Lucas Benedicic et al. 2019. Sarus: Highly scalable Docker containers for HPC systems. In *Proc. ISC Workshops*. https://doi.org/10.1007/978-3-030-34356-9_5
 - [8] David Bernstein. 2014. Containers and cloud: From LXC to Docker to Kubernetes. *IEEE Cloud Computing* (2014). <https://doi.org/10.1109/MCC.2014.51>
 - [9] Jonathan Calmels et al. 2021. ENROOT. <https://github.com/NVIDIA/enroot/blob/v3.3.0/README.md>
 - [10] Byron Cook. 2018. Formal reasoning about the security of Amazon Web Services. In *Proc. Computer Aided Verification*. https://doi.org/10.1007/978-3-319-96145-3_3
 - [11] J.H.M. Dassen, joost witteven, and Clint Adams. 2019. *faked(1)*. <https://manpages.debian.org/buster/fakeroot/faked.1.en.html>
 - [12] J.H.M. Dassen, joost witteven, and Clint Adams. 2019. *fakeroor(1)*. <https://manpages.debian.org/buster/fakeroot/fakeroor.1.en.html>
 - [13] Debian Project. 2021. *fakeroor* (GitLab repository). <https://salsa.debian.org/clint/fakeroot>
 - [14] Adam Di Carlo et al. 2019. *Release notes for Debian 9 (Stretch)*. Chapter 5.3.2.1. <https://www.debian.org/releases/stretch/amd64/release-notes/ch-information.en.html#apt-issues>
 - [15] Docker, Inc. 2021. Run the Docker daemon as a non-root user. <https://docs.docker.com/engine/security/rootless>
 - [16] Wes Felter, Alexandre Ferreira, Ram Rajamony, and Juan Rubio. 2015. An updated performance comparison of virtual machines and Linux containers. In *Proc. Performance Analysis of Systems and Software*. <https://doi.org/10.1109/ISPASS.2015.7095802>
 - [17] Kurt B. Ferreira, Patrick Bridges, and Ron Brightwell. 2008. Characterizing application sensitivity to OS interference using kernel-level noise injection. In *Proc. SC*.
 - [18] Todd Gamblin et al. 2015. The Spack package manager: Bringing order to HPC software chaos. In *Proc. SC*. <https://dl.acm.org/doi/10.1145/2807591.2807623>
 - [19] Holger Gantkow, Steffen Walter, and Christoph Reich. 2020. Rootless containers with Podman for HPC. In *Proc. ISC*. https://doi.org/10.1007/978-3-030-59851-8_23
 - [20] Lisa Gerhardt et al. 2017. Shifter: Containers for HPC. In *Journal of Physics: Conference Series*. <https://doi.org/10.1088/1742-6596/898/8/082021>
 - [21] Talal Halabi and Martine Bellaiche. 2017. Towards quantification and evaluation of security of cloud service providers. *Information Security and Applications* (2017). <https://doi.org/10.1016/j.jisa.2017.01.007>
 - [22] Amit Hendre and Karuna Pande Joshi. 2015. A semantic approach to cloud security and compliance. In *Proc. Cloud Computing*. <https://doi.org/10.1109/CLOUD.2015.157>
 - [23] William Henry. 2019. Podman and Buildah for Docker users. <https://developers.redhat.com/blog/2019/02/21/podman-and-buildah-for-docker-users/>
 - [24] Douglas M. Jacobsen and Richard Shane Canon. 2015. Contain this, unleashing Docker for HPC. In *Proc. Cray User Group*. <http://www.nersc.gov/assets/Uploads/cug2015udi.pdf>
 - [25] Jay Fenlason et al. John Gilmore. 2021. *GNU tar: An archiver tool* (1.34 ed.). <https://www.gnu.org/software/tar/manual/>
 - [26] Poul-Henning Kamp and Robert N.M. Watson. 2000. Jails: Confining the omnipotent root. In *Proc. SANE*. <http://www.sane.nl/events/sane2000/papers/kamp.pdf>
 - [27] Michael Kerrisk. 2013. Namespaces in operation, part 1: Namespaces overview. *Linux Weekly News* (Jan. 2013). <https://lwn.net/Articles/531114/>
 - [28] Michael Kerrisk. 2013. Namespaces in operation, part 5: User namespaces. *Linux Weekly News* (Feb. 2013). <https://lwn.net/Articles/532593/>
 - [29] Michael Kerrisk et al. 2017. *setresuid(2)* (5.10 ed.). <https://man7.org/linux/man-pages/man2/setresuid.2.html>
 - [30] Michael Kerrisk et al. 2019. *getgroups(2)* (5.10 ed.). <https://man7.org/linux/man-pages/man2/getgroups.2.html>
 - [31] Michael Kerrisk et al. 2020. *capabilities(7)* (5.10 ed.). <https://man7.org/linux/man-pages/man7/capabilities.7.html>
 - [32] Michael Kerrisk et al. 2020. *chown(2)* (5.10 ed.). <https://man7.org/linux/man-pages/man2/chown.2.html>
 - [33] Michael Kerrisk et al. 2020. *mount_namespaces(7)* (5.10 ed.). https://man7.org/linux/man-pages/man7/mount_namespaces.7.html
 - [34] Michael Kerrisk et al. 2020. *user_namespaces(7)* (5.10 ed.). https://man7.org/linux/man-pages/man7/user_namespaces.7.html
 - [35] Gregory M. Kurtzer, Vanessa Sochat, and Michael W. Bauer. 2017. Singularity: Scientific containers for mobility of compute. *PLOS ONE* (May 2017). <https://doi.org/10.1371/journal.pone.0177459>
 - [36] Emily Le and David Paz. 2017. Performance analysis of applications using Singularity container on SDCS Comet. In *Proc. PEARC*. <https://doi.org/10.1145/3093338.3106737>
 - [37] Peini Liu and Jordi Guitart. 2020. Performance comparison of multi-container deployment schemes for HPC workloads: an empirical study. *Supercomputing* (2020). <https://doi.org/10.1007/s11227-020-03518-1>
 - [38] Manoj Naik and Marc Eshel. 2017. *File system extended attributes in NFSv4*. Proposed standard RFC 8276. IETF. <https://datatracker.ietf.org/doc/rfc8276/>
 - [39] Claus Pahl and Pooyan Jamshidi. 2016. Microservices: A systematic mapping study. In *Proc. Cloud Computing and Services Science*. <https://dl.acm.org/doi/10.5220/0005785501370146>
 - [40] Kevin Pedretti et al. 2020. Chronicles of Astra: Challenges and lessons from the first petascale Arm supercomputer. In *Proc. SC*. <https://dl.acm.org/doi/abs/10.5555/3433701.3433764>
 - [41] Fabrizio Petrini, Darren J Kerbyson, and Scott Pakin. 2003. The case of the missing supercomputer performance: Achieving optimal performance on the 8,192 processors of ASCI Q. In *Proc. SC*. <https://doi.org/10.1145/1048935.1050204>
 - [42] Reid Priedhorsky and Tim Randles. 2017. Charliecloud: Unprivileged containers for user-defined software stacks in HPC. In *Proc. SC*. <https://doi.org/10.1145/3126908.3126925>
 - [43] Ling Qian et al. 2009. Cloud computing: An overview. In *Proc. Cloud Computing*. https://doi.org/10.1007/978-3-642-10665-1_63
 - [44] Ajeet Raina. 2018. Under the hood: Demystifying Docker for Mac CE Edition. <https://collabnix.com/how-docker-for-mac-works-under-the-hood>
 - [45] Red Hat, Inc. 2018. 7.6 release notes. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html-single/7.6_release_notes/index#appe-7.6_Release_Notes-Revision_History
 - [46] Reventlov. 2015. Using the docker command to root the host (totally not a security issue). <https://web.archive.org/web/20170614013206/http://reventlov.com/advisories/using-the-docker-command-to-root-the-host>
 - [47] Seeks. 2018. index: pseudo (GitLab repository). <http://git.yoctoproject.org/cgit/cgit.cgi/pseudo/tree/?h=pseudo-1.9.0>
 - [48] Shemesh Shachar. 2019. Fakeroor NG. https://fakeroor-ng.lingnu.com/index.php?title=Main_Page&oldid=2
 - [49] shadow-utils project. 2020. *newgidmap(1)* (4.8.1 ed.). <https://man7.org/linux/man-pages/man1/newgidmap.1.html>
 - [50] shadow-utils project. 2020. *newuidmap(1)* (4.8.1 ed.). <https://man7.org/linux/man-pages/man1/newuidmap.1.html>
 - [51] Schacher Shemesh. 2016. Fakeroor Next Gen (SourceForge repository). <https://sourceforge.net/projects/fakeroorng>
 - [52] Jonathan Sparks. 2017. HPC containers in use. In *Proc. Cray User Group*. https://cug.org/proceedings/cug2017_proceedings/includes/files/pap164s2-file1.pdf
 - [53] Sylabs, Inc. 2021. Singularity Enterprise. <https://sylabs.io/singularity-enterprise>
 - [54] Alfred Torrez, Timothy Randles, and Reid Priedhorsky. 2019. HPC container runtimes have minimal or no performance impact. In *Proc. CANOPIE-HPC*. <https://conferences.computer.org/sc19w/2019/pdfs/CANOPIE-HPC2019-7nd7J7oXB1gtzeJIH79mM/4VQ694HccmYWPtiazfngGiz/RirsVUUC5jhKx5UKJB1yT.pdf>
 - [55] Linus Torvalds. 2013. Linux 3.8. *linux-kernel mailing list* (Feb. 2013). <http://lkml.iu.edu/hypermail/linux/kernel/1302.2/00936.html>
 - [56] Dan Walsh. 2018. Reintroduction of Podman. <https://www.projectatomic.io/blog/2018/02/reintroduction-podman/>
 - [57] Dan Walsh et al. 2021. Basic setup and use of Podman in a rootless environment. <https://github.com/containers/podman/blob/master/docs/tutorials/rootless-tutorial.md>
 - [58] Lizhe Wang et al. 2010. Cloud computing: A perspective study. *New Generation Computing* (2010). <https://doi.org/10.1007/s00354-008-0081-5>
 - [59] Peter Wienemann. 2021. Personal communication.
 - [60] Miguel G. Xavier et al. 2013. Performance evaluation of container-based virtualization for high performance computing environments. In *Proc. Parallel, Distributed, and Network-Based Processing*. <https://doi.org/10.1109/PDP.2013.41>
 - [61] Andrew J. Younge et al. 2017. A tale of two systems: Using containers to deploy HPC applications on supercomputers and clouds. In *Proc. CloudCom*. <https://doi.org/10.1109/CloudCom.2017.40>
 - [62] Andrew J. Younge et al. 2020. Advanced tri-lab software environment (ATSE), version 1.2.5. <https://doi.org/10.11578/dc.20201109.3>

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

Content for the figures in the paper was generated with a simple Makefile-based procedure, based on processing Dockerfiles with Charliecloud and Podman, along with some stand-alone shell scripts. Details are in the provided artifact archive's README. The archive contains all necessary code. See this artifact for more environment information details.

The second artifact listed is the Charliecloud source code, containing the new features described in the paper.

Author-Created or Modified Artifacts:

Persistent ID:

↪ <https://doi.org/10.6084/m9.figshare.14396099>

Artifact name: Examples

Persistent ID: <https://github.com/hpc/charliecloud/releases/tag/v0.23>

↪ <https://github.com/hpc/charliecloud/releases/tag/v0.23>

Artifact name: Charliecloud source code

Citation of artifact:

↪ <https://doi.org/10.1145/3126908.3126925>

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: Generic x86 VM; Astra ARM super-computer

Operating systems and versions: Debian 10; RHEL7 derivative

Applications and versions: Charliecloud 0.23; Podman 1.4.4 and 3.0.1

Input datasets and versions: Dockerfiles included