

Void Processes: Minimising privilege by default on Linux

Jake Hillion
University of Cambridge
United Kingdom
jake.hillion@cl.cam.ac.uk

Abstract

Operating systems are providing more facilities for process isolation than ever before, realised in technologies such as Containers [CN] and systemd slices [CN]. These systems separate the design of the program from the systems that create privilege separation.

Void Processes take these techniques to the extreme, removing access to everything but syscalls from a process by default. This work focuses on adding back slivers of privilege to achieve functional applications with minimal privilege.

I present a summary of the privilege separation features in modern Linux, the system design of void processes, the language front-ends to support it, and an evaluation on a series of example applications.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

Keywords: datasets, neural networks, gaze detection, text tagging

ACM Reference Format:

Jake Hillion. 2022. Void Processes: Minimising privilege by default on Linux. *J. ACM* 37, 4, Article 111 (August 2022), 8 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

Void processes take advantage of modern Linux namespaces to attempt to run applications without exposing them to the system itself. Void processes use a mixture of Linux namespaces and file descriptive based capabilities to allow running purpose-built applications without expecting the support of the standard Linux system. During the process of building

Author's address: Jake Hillion, University of Cambridge, United Kingdom, jake.hillion@cl.cam.ac.uk.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

0004-5411/2022/8-ART111 \$15.00

<https://doi.org/XXXXXXXX.XXXXXXX>

such a system, gaps in the kernel were exposed, given that this work is at the edge of what main spaces can achieve. This work will go onto detail the process of creating void processes themselves, re-adding features that these processes need to do useful work, and the learnings of what features are missing in the user-space kernel APIs to succeed in creating processes this way.

This work explores the question of what is an operating system by taking a novel approach to running applications with the system exposed in an entirely different way. Rather than limiting the access of a process or set of processes to the operating system, such as in containers, we instead limit the access to the operating system with more explicit methods per process. Interaction between processes is allowed by specifying such interaction statically at compile time, removing any separation between the application developer and the system controlling access to the application, unlike solutions such as SELinux.

2 Motivation

This work aims to explore the limits of the Linux userspace APIs in the context of complete process isolation, producing a software ecosystem to support running applications with fully minimised privilege. Further, discussion will be made of which parts of the API are well-suited and which are not, and how they might be better designed. Finally, the performance of absolute separation is evaluated, to find out at what cost this can be achieved in the current kernel.

Void processes aim to serve a different purpose than containers. Rather than virtualising a different operating system with the same kernel, void processes aim to remove all but what is necessary to run the application from the current operating system. The cut down appearance of the OS is still the same OS, rather than bringing in new utilities and libraries from perhaps a different OS. It would be possible to include this feature as an extension, but it is not the primary goal. This difference is important when deciding the features of the application, particularly the decision to exclude time namespaces (§4.1.8).

2.1 Threat Model

I present a threat model in which application binaries are trusted absolutely. That is, the software provider had no ill intent, and once the binary is on disk, it will not change

without permission. This means that one can trust the binary to set up its own security, as it is protecting not against malice by its own developers, but instead bugs in the software.

Expand and finalise threat model

3 Background

3.1 Mount Namespaces

Mount namespaces were by far the most challenging part of this project. When adding new features, they continuously raised problems in both API description, expected behaviour, and performance of the tools given. A comparison will be given in this section to two other namespaces, network and UTS, to show the significant differences in the design goals of mount namespaces. Much of the programming issue here comes from a fundamental lack of consistency between mount namespaces and other namespaces in Linux, which will be discussed further in this section.

3.1.1 Copy-on-Write. Comparing to network namespaces, a slightly more modern namespace [Table 1], we see a huge difference in what occurs when a new namespace is created. When creating a new network namespace, one is immediately placed into a void, a network namespace containing only a loopback adapter. That is, the process has no ability to interact with the outside network, and no immediate relation to the parent network namespace. To interact with alternate namespaces, one must explicitly create a connection between the two, or move a physical adapter into the new (empty) namespace. Mount namespaces, rather than creating a new and empty namespace, made the choice to create a copy of the parent namespace, in a copy-on-write fashion. That is, after creating a new mount namespace, the mount hierarchy appears much the same as before.

3.1.2 Shared Subtrees. While some other namespaces are copy-on-write, for example UTS namespaces, they do not present the same problem as mount namespaces. Although UTS namespaces are copy-on-write, it is trivial to create a void by setting the hostname of the machine to a constant. This removes any relation to the parent namespace and to the outside machine. Mount namespaces instead maintain a shared pointer with most filesystems, more akin to not creating a new namespace than a copy-on-write namespace.

Shared subtrees [18] were introduced to provide a consistent view of the unified hierarchy between namespaces. Consider the example in Figure 1. `unshare(1)` creates a non-shared tree, which presents the behaviour shown. Although `/mnt/cdrom` from the parent namespace has been bind mounted in the new namespace, the content of `/mnt/cdrom` is not the same. This is because the filesystem newly mounted on `/mnt/cdrom` is unavailable in the separate mount namespace. To combat this, shared subtrees were introduced. That is, as long as `/mnt/cdrom` resides on a shared subtree, the

newly mounted filesystem will be available to a copy of `/mnt/cdrom` in another namespace.

`systemd` made the choice to mount `/` as a shared subtree [11]:

“Notwithstanding the fact that the default propagation type for new mount is in many cases `MS_PRIVATE`, `MS_SHARED` is typically more useful. For this reason, `systemd(1)` automatically remounts all mounts as `MS_SHARED` on system startup. Thus, on most modern systems, the default propagation type is in practice `MS_SHARED`.”

This means that when creating a new namespace, mounts and unmounts are propagated in by default. Further, it means that mounts and unmounts are propagated out of the namespace. This can be highly confusing behaviour, and `unshare(1)` considers this behaviour inconsistent with the goals of unsharing - it immediately calls `mount(none , / , NULL , MS_REC | MS_PRIVATE , NULL)` after `unshare(CLONE_NEWNS)`. The reasoning for this is that containers created should not present the behaviour given in Figure 1, and this behaviour is unavoidable unless the parent mounts are shared, while it is possible to disable the behaviour where necessary.

3.1.3 Lazy unmounting. Mount namespaces present further interesting behaviour when unmounting initial root filesystem. Although this may initially seem isolated to void processes, it is also a problem in a container type system. Consider again the container created in Figure 1 - the existing root must be unmounted after pivoting, to avoid keeping the container fully connected to the outside root.

Referring again to network namespaces, sockets continue to exist in their initial namespace, allowing for regular file-descriptor passing semantics [8]. Extending upon this socket behaviour is Wireguard, which creates adapters that may be freely moved between namespaces while continuing to connect externally from their initial parent [9, §7.3].

Something which behaves differently is the memory mapping of a currently running process's binary. Considering the example in Listing 1, which shows a short C program and the result of running it, it is seen that the `/` mount is busy when attempting the unmount. Given that the process was created in the parent namespace, the behaviour of file descriptors would suggest that the process would maintain a link to the parent namespace for its own memory mapped regions. However, the fact that the otherwise empty namespace has a busy mount shows that this is not the case.

A feature called lazy unmounting or `MNT_DETACH` exists for situations where a busy mount still needs to be unmounted. Supplying the `MNT_DETACH` flag to `umount2(2)` causes the mount to be immediately detached from the unified hierarchy, while remaining mounted until the last user has finished with it. While this initially seems like a good solution, this syscall is incredibly dangerous when combined with shared subtrees. This behaviour is shown in Figure 2, where a lazy

```
#
#
#
#
#
# mount /dev/sr0 /mnt/cdrom
# ls /mnt/cdrom
file_1 file_2
```

```
# unshare -m
# mount_container_root /tmp/a
# mount --bind \
    /mnt/cdrom /tmp/a/mnt/cdrom
# pivot_root /tmp/a /tmp/a/oldroot
# umount /tmp/a/oldroot
#
# ls /mnt/cdrom
```

Figure 1. Highly separated behaviour without shared subtrees between mount namespaces.

Listing 1. Behaviour when attempting to unmount / after an unshare.

```
int main() {
    if (unshare(CLONE_NEWNS))
        perror("unshare");
    if (mount("none", "/", NULL,
            MS_REC|MS_PRIVATE, NULL))
        perror("mount");
    if (umount("/"))
        perror("umount");
}
--
umount: Device or resource busy
```

(and hence recursive) unmount is combined with a shared subtree to disastrous effect.

This behaviour raises questions about why a shared subtree, which exists as an object, would need to be detached recursively - decreasing the reference count to the shared subtree itself would seem sufficient. The inconsistency is best explained by looking at the development timeline for the three features here: mount namespaces, shared subtrees, and recursive lazy unmounts.

When lazy unmounting was added, in September 2001, the author said the following [23]:

“There are only two things to take care of - a) if we detach a parent we should do it for all children b) we should not mount anything on “floating” vfstmounts. Both are obviously satisfied for current code (presence of children means that vfstmount is busy and we can’t mount on something that doesn’t exist).”

This logic held even in the presence of namespaces, with the initial patchset in February 2001 [23], as mounts were not initially shared but duplicated between namespaces. However, when shared subtrees were added in January 2005 [25], this logic stopped holding.

When setting up a container environment, one calls `pivot_root(2)` to replace the old root with a new root for the container.

Then, the old root may be unmounted. Oftentimes the solution is to exec a binary in the new root first, meaning that the old root is no longer in use and may be unmounted. This works, as old root is only a reference in this namespace, and hence may be unmounted with children - the vfstmount in this namespace is not busy, in contradiction to the quotation.

If, instead, one wishes to continue running the existing binary, this is possible with lazy unmounting. However, the kernel only exposes a recursive lazy unmount. With shared subtrees, this results in destroying the parent tree. While this is avoidable by removing the shared propagation from the subtree before unmounting, the choice to have `MNT_DETACH` aggressively cross shared subtrees can be highly confusing, and perhaps undesired behaviour in a world with shared subtrees.

4 System Design

An example of running a multi-entripoint application is given in Figure 3. What was originally a monolithic application becomes a set of applications that communicate with a new shim. The shim does not replace the kernel, and instead supplements it with new higher-level abilities. Each entripoint receives input from the shim, and can return data to the shim where appropriate. Most of this data is in the form of file descriptors, which are treated as capabilities in this system.

A multi-entripoint application stores the requirements for running it as static data in the ELF of the binary. When launched, `binfmt_misc` is used to launch the application with the multi-entripoint shim. The shim decodes this data and sets up processes and IPC accordingly.

The shim takes advantage of high levels of privilege to be able to more effectively deprive an application than an application with ambient authority could. For example, creating a new network namespace requires `CAP_SYS_ADMIN`, which would give many applications more privilege than they require. By deferring to a shim with extra privileges, this trusted code can be written only once, and avoid conferring more privileges than otherwise required.

```
# cat /proc/mounts | grep udev
udev /dev devtmpfs rw,nosuid,relati...
#
#
# cat /proc/mounts | grep udev
cat: /proc/mounts: No such file or ...
```

```
#
#
# unshare --propagation unchanged -m
# umount -l /
#
#
```

Figure 2. Behaviour when attempting to unmount / from an unshared shell with a shared mount.




Figure 3. Interaction between the application and the environment.

4.1 Building the Void

Preparing a void for a new process takes advantage of the namespaces feature in Linux. However, many of the namespaces are not designed for this purpose, so this is a more difficult prospect than one might hope. Details of when each namespace was added and some of the relevant features are given in Table 1.

4.1.1 Mount namespaces. Mount namespaces were the first [Table 1] namespaces introduced to Linux, in kernel version 2.5.2 [CN]. In contrast to network namespaces, the API is particularly unfriendly to creating a Void. The creation of mount namespaces is copy-on-write, and many filesystems are mounted shared. This means that they propagate changes back through namespace boundaries. As the mount namespace does not allow for creating an entirely new root, extra care must be taken in separating processes. The method taken in this system is mounting a new tmpfs file system in a new namespace, which doesn't propagate to the parent, and using the pivot_root(8) command to make this the new root. By pivoting to the tmpfs without bind mounting the old root inside, the old root becomes completely inaccessible from the namespace. Similarly, the tmpfs never appears in the parent namespace. Finally, after ensuring the old root is set to MNT_PRIVATE to avoid propagation (more details in §3.1.2), the old root can be lazily detached. This allows the binary from the parent namespace, the shim in this case, to continue running correctly. Any new processes only have access to the materials in the tmpfs void.

4.1.2 IPC namespaces. Creating a void with IPC namespaces is pleasantly easy in comparison. From the manual page [10]:

“Objects created in an IPC namespace are visible to all other processes that are members of that namespace, but are not visible to processes in other IPC namespaces.”

This provides exactly the correct semantics for a void, particularly because it is not copy-on-write. IPC objects are visible within a namespace if and only if they are created within that namespace.

4.1.3 UTS namespaces. UTS namespaces provide isolation of the hostname and domain name of a system between processes. Similarly to IPC namespaces, all processes in the same namespace see the same results for each of these. Unlike IPC namespaces, UTS namespaces are copy-on-write. That is, the value of each of these in the parent namespace is the same in the child.

As the copied value does give information about the world outside of the void process, slightly more must be done than placing the process in a new namespace. Fortunately this is easy for UTS namespaces, as the hostname and domain name can be set to constants, removing any link to the parent.

4.1.4 user namespaces.

```
user namespaces
```

4.1.5 Network namespaces. Network namespaces were added in kernel version 2.6.24 [3], some time after the initial namespace boom. They present the optimal namespace for creating a void. Creating a new network namespace immediately creates an entirely empty namespace. That is, the new network namespace has no link whatsoever to the creating network namespace. To add a link, one can create a virtual Ethernet pair, with one adapter in each namespace [CN]. Alternatively, one can create a Wireguard adapter with sending and receiving sockets in one namespace and the VPN adapter in another [9, §7.3]. This allows for very high levels of separation while still maintaining access to the primary resource - the Internet or wider network.

4.1.6 PID namespaces. pid namespaces add a mapping from the process IDs inside the namespace to process IDs in the parent namespace. This continues until processes reach

Table 1. Table showing the date and kernel version each namespace was added. The date provides the first commit where they appeared date of creation, and the kernel version the kernel release they appear in the changelog of. Namespaces are ordered by kernel version then alphabetically. A checkbox is provided for whether the namespace displays copy-on-write behaviour, discussed in more detail in Section 4.1.

Namespace	CoW	Date	Kernel Version
mount	^a	24th February 2001 [24]	2.5.2 [20]
ipc		2nd October 2006 [16]	2.6.19 [1]
uts		2nd October 2006 [14]	2.6.19 [1]
user		15th July 2007 [17]	2.6.23 [2]
network		10th October 2007 [7]	2.6.24 [3]
pid		2nd October 2006 [6]	2.6.24 [3]
cgroup	^b	18th March 2016 [15]	4.6 [21]
time		12th November 2019 [22]	5.6 [4]

^aShared namespaces make for unique behaviour, see §3.1.2.

^bcgroup namespaces provide a virtualised view of the unified cgroup namespace and do not allow private modification.

the top-level pid namespace. This isolation behaviour is different to that of some other namespaces, as each process within the namespace represents a process in the parent namespace too.

Although pid namespaces work quite well for creating a void from the perspective of the inside process, some care must be taken in the implementation, as the actions of pid namespaces are highly affected by others. Some examples of this slightly unusual behaviour are shown in Listing 2.

The first behaviour shown is that a `unshare(CLONE_PID)` call followed immediately by `execd` does not have the desired behaviour. The reason for this is that the first process created in the new namespace is given PID 1 and acts as an init process. That is, whichever process the shell spawns first becomes the init process of the namespace, and when that process dies, the namespace can no longer create new processes. This behaviour is avoided by either calling `unshare/fork`, or utilising `clone(2)` instead. The `unshare(3)` binary provides a `fork` flag to solve this, while the implementation of the void orchestrator uses `clone(2)` which combines the two into a single syscall.

Secondly, we see that even in a shell that appears to be working correctly, processes from outside of the new pid namespace are still visible. This behaviour occurs because the mount of `/proc` visible to the process in the new pid namespace is the same as the init process. This is solved by remounting `/proc`, available through `unshare(3)` with the `mount-proc` flag. Care must be taken that this mount is completed in a new mount namespace, or else processes outside of the pid namespace will be affected. The void orchestrator again avoids this by voiding the mount namespace entirely, so any access to `proc` must be either bound to outside the namespace, or freshly mounted.

4.1.7 cgroup namespaces.

cgroup namespaces

Listing 2. Unshare behaviour with pid namespaces, with and without forking and remounting `proc`.

```

$ unshare =p
=bash: fork: Cannot allocate memory
# (new shell in new pid namespace)
# ps ax | tail =n 3
=bash: fork: Cannot allocate memory

$ unshare ==fork =p
# (new shell in new pid namespace)
# ps ax | tail =n 3
2645 ?          I          0:00 [kworker/...]
2689 pts/1      R+         0:00 ps ax
2690 pts/1      S+         0:00 tail =n 2

$ unshare ==fork ==mount =proc =p
# (new shell in new pid namespace)
# ps ax | tail =n 3
1 pts/1        S          0:00 =bash
15 pts/1       R+         0:00 ps ax
16 pts/1       S+         0:00 tail =n 3
    
```

4.1.8 time namespaces. Time namespaces are the final namespace added at the time of writing, added in kernel version 5.6 [4]. The motivation for adding time namespaces is given in the manual page [12]:

The motivation for adding time namespaces was to allow the monotonic and boot-time clocks to maintain consistent values during container migration and checkpoint/restore.

That is, time namespaces virtualise the appearance of system uptime to processes, rather than attempting to virtualise the wall clock time. This is important for processes that depend on it in one specific situation: migration. If an uptime

dependent process is migrated from a machine that has been up for a week to a machine that was booted a minute ago, the guarantees provided by the `CLOCK_MONOTONIC` and `CLONE_BOOTTIME` no longer hold.

This results in time namespaces having very limited usefulness in a system that does not support migration, such as the one presented here. Perhaps randomised o sets would hide some information about the system, but the usefulness is debatable and the quantity of bespoke syscalls would slow down the application. Time namespaces are thus avoided in this implementation.

4.2 Filling the void

Once a void has been created the goal is to reinsert enough to run the application, and no more. To allow for running applications in the void with minimal kernel changes, this is done using a mixture of `le`-descriptor capabilities and adding elements to the namespaces. Capabilities allow for a clean experience where suitable, while adding elements to namespaces creates a more Linux-like experience for the application.

4.2.1 Files and directories. There are two options to provide access to `les` and directories in the void. Firstly, for a single `le`, an already open `le` descriptor can be o ered. Consider the TLS broker of a TLS server with a persistent certificate and key `le`. Only these `les` are required to correctly run the application - no view of a `lesystem` is necessary. Providing an already opened `le` descriptor gives the process a capability to those `les` while requiring no concept of a `lesystem`, allowing that to remain a complete void. This is possible because of the semantics of `le` descriptor passing across namespaces - the `le` descriptor remains a capability, regardless of moving into a namespace without access to the `le` in question.

Alternatively, `les` and directories can be mounted in the void namespace. This supports three things which the capabilities do not: directories, dynamic linking, and applications which have not been adapted to use `le` descriptors. Firstly, the existing `openat(2)` calls are not suitable by default to treat directory `le` descriptors as capabilities, as they allow the search path to be absolute. This means that a process with a directory `le` descriptor in another namespace can access any `les` in that namespace [RN] by supplying an absolute path. Secondly, dynamic linking is best served by binding `les`, as these read only copies and the trusted binaries ensure that only the required libraries can be linked against. Finally, support for individual required `les` can be added by using `le` descriptors, but many applications will not trivially support it. Binding `les` allows for a form of backwards compatibility.

4.2.2 Networking.

Write about reintroducing networking to the void.

Listing 3. A sample application using the Rust language frontend.

```
#[entrypoint]
fn encrypt(mut in: File , mut out: File )

#[entrypoint]
fn main() {
    let input_file = ...;
    let output_file = ...;

    encrypt(input_file , output_file );
}
```

5 Language Frontends

The language frontends are an extremely important part of this project, closing the gap between a static privilege separation solution like SELinux [CN] and a dynamic one like Capsicum [7]. I have implemented a language frontend in Rust and will describe it in this section.

5.1 Rust

The Rust frontend uses macros to wrap functions with high-level primitives into multi-entrypoint compatible entrypoints. Further, it allows calling these functions using the new interface via the shim. Consider the example in Listing 3.

Firstly, the `encrypt` entrypoint is created. This is a regular Rust function which takes two high-level `File` objects, a wrapped `le` descriptor. The entrypoint macro wraps this function, providing in its place a `extern "C"` function that is unmangled and takes `argc/argv`. This allows functions with high-level arguments to be used as normal, with the argument parsing abstracted away by the library.

Second is the ordinary `main` function for the application. This is also tagged as an entrypoint, allowing the library to help out with more calls. The example given here is that of the `encrypt` method, which uses the API seen above. The use of macros here allows the call to `encrypt` to remain type safe, even though the call must pass through an external interface (the shim itself).

A significant benefit to this approach is the ease of disabling the multi-entrypoint application. By turning the entrypoint macro into identity with a `crate` feature, the code is compiled without the aid of the multi-entrypoint shim. This allows for significantly easier debugging, as the application follows a single execution path, rather than needing to be debugged as a distributed application.

Listing 4. An application that requires only stdout and stderr.

```
#[entrypoint(stdout)]
fn main() { println!("hello world!"); }
```

Figure 4. Process separation in a TLS server.

6 Example Applications

6.1 No Permissions

The cornerstone of strong process separation is an application that is completely depriveged. Listing 4 shows an application which, when run under the shim, drops all privileges except stdout. This is easy to achieve under the shim.

6.2 gzip

GNU gzip [13] is well structured for privilege separation, though doesn't implement it by default. There is a clear split between the processing logic, selecting the items to do work on, and the compression/decompression routines, each of which are handed a pair of input and output file descriptors. This is shown by Watson et al. in [27].

As C does not have high-level language features for multi-entrypoint applications, adapting it is slightly more verbose than the other examples seen. However, the resulting code change is still only X lines, if a bit more intricate. This places the risky compression and decompression routines in full sandboxes, while still allowing the simpler argument processing code ambient authority. The argument processing code needs no additional Linux capabilities to manage this permissioning, as the required capabilities are provided by the shim.

6.3 TLS Server

Finally, a rudimentary TLS server is created to show the rich privilege separation abilities of multi-entrypoint applications. An example structure is shown in Figure 4. Rather than being provided with a view of the network, the initial TCP handling

process is given an already bound socket listener by the shim. This allows the TCP handler to live in an extremely restricted zero-access network namespace, while still performing the tasks of receiving new TCP connections.

Next, the TCP handler hands off the new TCP connections to the shim. Though the figure shows this as a direct connection between the TCP handler and the TLS handler, they are passed through the shim, from which the shim spawns a fresh TLS handler for each connection. The TLS handler is handed file descriptors to the certificate and keys that it requires, and hands back a decrypted request reader and an empty response writer file descriptor to the shim.

Finally, this pair of decrypted request reader and response writer are handed to a new process which handles the request. In the example case, this new process is handed a dirfd to /var/www/html, which is bind-mounted into an empty file system namespace by the shim. This allows the request handler enough access to serve files, while restricting access to anything else.

7 Evaluation

Write evaluation

8 Related Work

8.1 Virtual Machines and Containers

Virtual Machine solutions [5, 26] provide the ability to split a single machine into multiple virtual machines. When placing a single application in each virtual machine, they are effectively isolated from one another. Full fat container solutions such as Docker [CN], containerd [CN], and systemd-ns [CN] provide mechanisms to isolate an application almost completely from other applications running on a single machine. Some have claimed that this provides isolation superior to virtual machines [19].

Both of these solutions are less effective at isolating parts of an application from itself [CN with research]. Consider running only a TLS web server in a virtual machine. Although other applications will be unable to access the certificates, as they are in different virtual machines, methods within the application that should not be able to access the certificates still can.

While virtual machines and containers provide a strong isolation at the application level, they are not a compelling solution to intra-application privilege separation.

8.2 systemd

systemd [CN] provides a declarative interface to all of the process separation techniques used in this work. Rather than the responsibility of the programmer, creating these declarative descriptions is most commonly left to the package maintainers. This work seeks to provide similar capabilities to the people best suited to privilege separating an application: the developers.

8.3 Capsicum

Capsicum [27] extends UNIX file descriptors in FreeBSD to reflect the rights on the object they hold. These capabilities may be shared between processes as other file descriptors. The goals of both software are the same: make privilege separated software better. However, we take quite different approaches. Multi-entrypoint applications focus on building a static definition really close to the code, while Capsicum allows processes to dynamically privilege separate. This allows applying static analysis to the policies, while also keeping the definition close to the code.

9 Future Work

9.1 Dynamic Linking

Dynamic linking works correctly under the shim, however, it currently requires a high level of manual input. Given that the threat model in Section 2.1 specifies trusted binaries, it is feasible to add a pre-spawning phase which appends read-only libraries to the specification for each spawned process automatically before creating appropriate voids. This would allow anything which can link correctly on the host system to link correctly in void processes.

10 Conclusion

Write conclusion

Acknowledgments

Write acknowledgements

References

- [1] 2006. Linux Version 2.6.19 Changelog. https://kernelnewbies.org/Linux_2_6_19
- [2] 2007. Linux Version 2.6.23 Changelog. https://kernelnewbies.org/Linux_2_6_23
- [3] 2008. Linux Version 2.6.24 Changelog. https://kernelnewbies.org/Linux_2_6_24
- [4] 2020. Linux Version 5.6 Changelog. https://kernelnewbies.org/Linux_5.6
- [5] Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. 2003. Xen and the art of virtualization. *ACM SIGOPS Operating Systems Review* 37, 5 (Oct. 2003), 164–177. <https://doi.org/10.1145/1165389.945462>
- [6] Sukadev Bhattiprolu. 2006. [PATCH] Define struct pspage. <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=3fbc96486459324e182717b03c50c90c880be6ec>
- [7] Eric W. Biederman. 2007. [NET]: Basic network namespace infrastructure. <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=5f256becd868bf63b70da8f2769033d6734670e9>
- [8] Eric W. Biederman. 2007. Re: netns: close all sockets at unshare? <https://lore.kernel.org/all/m1r6kcccxw.fsf@ebiederm.dsl.xmission.com/>
- [9] Jason A. Donenfeld. 2017. WireGuard: Next Generation Kernel Network Tunnel. In *Proceedings 2017 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. <https://doi.org/10.14722/ndss.2017.23160>
- [10] Free Software Foundation. 2021. ipc_namespaces(7). https://man7.org/linux/man-pages/man7/ipc_namespaces.7.html
- [11] Free Software Foundation. 2021. mount_namespaces(7). https://man7.org/linux/man-pages/man7/mount_namespaces.7.html
- [12] Free Software Foundation. 2021. time_namespaces(7). https://man7.org/linux/man-pages/man7/time_namespaces.7.html
- [13] Jean-loup Gailly. 2020. Gzip. <https://www.gnu.org/software/gzip/>
- [14] Serge E. Hallyn. 2006. [PATCH] namespaces: utsname: implement utsname namespaces. <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=4865ecf1315b450ab3317a745a6678c04d311e40>
- [15] Tejun Heo. 2016. [GIT PULL] cgroup namespace support for v4.6-rc1. <https://lore.kernel.org/all/20160318190919.GF20028@mtj.duckdns.org/#r>
- [16] Kirill Korotaev. 2006. [PATCH] IPC namespace core. <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=25b21cb2f6d69b0475b134e0a3e8e269137270fa>
- [17] Cedric Le Goater. 2007. user namespace: add the framework. <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=acce292c82d4d82d3553b928df2b0597c3a9c78>
- [18] Ram Pai and Alexander Viro. 2005. Shared Subtrees. <https://www.kernel.org/doc/Documentation/filesystems/sharesubtree.txt> Added in commit 9cfceea8f7e8f5554e9c8130e568bcfa98a3a64.
- [19] Stephen Poltesz, Herbert Pözl, Marc E. Fiuczynski, Andy Bavier, and Larry Peterson. 2007. Container-based operating system virtualization: a scalable, high-performance alternative to hypervisors. In *Proceedings of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems 2007 (EuroSys '07)*. Association for Computing Machinery, New York, NY, USA, 275–287. <https://doi.org/10.1145/1272996.1273025>
- [20] Linus Torvalds. 2002. Linux Kernel Version 2.5.2 Changelog. <https://mirrors.edge.kernel.org/pub/linux/kernel/v2.5/ChangeLog-2.5.2>
- [21] Linus Torvalds. 2016. Linux 4.6-rc1. <https://lore.kernel.org/all/CA+55aFzBncC+F8TEb5KU1QVwA=PCA89mDp1VLNq008oZq8vpJQ@mail.gmail.com/>
- [22] Andrei Vagin. 2020. ns: Introduce Time Namespace. <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=769071ac9f20b6a447410c7eaa55d1a5233ef40c>
- [23] Alexander Viro. 2001. [PATCH] lazy umount (1/4). <https://lore.kernel.org/all/Pine.GSO.4.21.0109141427070.11172-100000@weyl.math.psu.edu/>
- [24] Alexander Viro. 2001. [PATCH][CFT] per-process namespaces for Linux. <https://lore.kernel.org/all/Pine.GSO.4.21.0102242253460.24312-100000@weyl.math.psu.edu/>
- [25] Alexander Viro. 2005. [RFC] shared subtrees. <https://lore.kernel.org/all/20050113221851.G126051@parcellfarce.linux.theplanet.co.uk/>
- [26] Inc. VMware. 2008. *Understanding Full Virtualization, Paravirtualization and Hardware Assist*. Technical Report. https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/VMware_paravirtualization.pdf
- [27] Robert NM Watson, Jonathan Anderson, Ben Laurie, and Kris Kennaway. 2010. Capsicum: Practical Capabilities for UNIX.. In *UNIX Security Symposium*, Vol. 46. 2.