It's Containers All The Way Down

by Aaron Grothe
www.grothe.us
June 15th, 2022

Introduction

Containers?

Containers have been around for over 40 years. Very timely topic :-)

The real change is how much easier it is to use containers nowadays. The first real step along this when LXC (LinuX Containers) made it into the Linux kernel, and then it took off when Docker came along.

Docker is designed pretty much for doing server stuff. You can run GUI stuff, but it takes a bit of work. Now tools such as distrobox have come along and are making it a lot easier.

Introduction (Continued)

If you have questions/comments please feel free to ask them anytime. You don't have to hold them until the end of the talk.

If there are other resources similar to these that you think might be useful to people please let the group know.

Hopefully this will be an interactive and productive session.

The slides for this are at the NEbraskaCERT website https://www.nebraskacert.org/csf and on my home page https://www.grothe.us/certifications

Demo

Want to do it a bit differently for this talk. So let's start with the demo.

My machine is a Debian box running Debian Unstable. Debian Unstable is the rolling release version of Debian. It is the latest and greatest and may or may not work.

We're going to use toolbox for this demonstration.

Demo

Let's confirm that we're running debian

% cat /etc/debian_version

Let's confirm that we're not running Redhat

% ls /etc/redhat-release

So we're running on a Debian Box with bookworm/sid

Demo

Let's start up the Fedora 36 environment

% toolbox enter

Does the system really think it is Fedora 36?

% cat /etc/redhat-release

Let's confirm that it doesn't think it is Debian

% ls /etc/debian_version

So we're on a Fedora 36 system and not a Debian System

Demo

Let's update the Fedora 36 environment

% yum upgrade

Fails we need to be root

% sudo yum upgrade

So I as a non-root user can do root things? Yep, that is how it is setup.

Demo

Let's fire up Firefox and see how that goes

% firefox

One last thing before we conclude the demo

% cd
% ls

It mounts my home directory into the environment, so all my files/data are available.

Demo

So that is a quick demo of what toolbox can do for us

Does a lot for us behind the scenes.

- Fires up podman
- Mounts the home directory
- Mounts proc, X11 and so on behind the scenes - so graphics, sound, etc all work
- Allows us to run as root inside the container (limited to what my userid can do)

Demo.

That is just a taste of what toolbox can do.

Toolbox can run a lot of different distributions

- Arch, Debian, Fedora, CentOS, SUSE and a lot of others. It can actually access about any container that is available, but that is a few of them

In the old days getting all this setup would take a lot of work and be beyond the average user. Nowadays, not a big deal.

Demo.

Toolbox was created by Redhat. It has several potential uses for it.

- It can be used on immutable systems such as SilverBlue to allow people to make changes
- Allows non-root users to be able to install tools/systems and run them
- Keeps the underlying operating system clean

Other Tools We're going to Talk about

Here are a list of the other tools we're going to talk about today

- KASM workspaces
- Distrobox
- Firejail
- Firefox Containers
  - Multi-User Account
  - Facebook Container
  - Google Containers
  - Many Others
- Sandboxie - have to show some love for the windows people :-)
- Containers you're not probably aware of

KASM Workspaces

KASM Workspaces is a very nice wrapper written around Docker and containers that make it possible for you to easily create/destroy containers.

Let's do another demo of KASM

Off to the browser we go

http://kasm.com

Let's fire up container and give it a run :-)

KASM Workspaces

Let's do something else interesting.

In my google browser on the main box I have installed the KASM extension for chrome. This will create a new window for a link and open it in a virtual docker instance on the remote box. It will also handle destroying the container when I'm done with it.

What does this do?

It gives me a safe place to open links I don't trust.

KASM Workspaces.

KASM is Open Source, but you'll probably want to get a support contract if you're using it.

The ability to quickly create remote images, safely open links and so on are pretty cool.

DistroBox

Let's take a quick peek at Distrobox

Distrobox should look pretty similar to toolbox for us

Let's start with listing the available images on the machine

% distrobox ls

To mix it up we'll do Almalinux this time round

% distrobox enter test

DistroBox

Should look a lot like toolbox

% cat /etc/redhat-release

We'll fire up firefox on this box as well

% firefox

Checking the about it is the AlmaLinux version of firefox

FireJail

FireJail uses Seccomp, cgroups and Linux namespaces to lock down applications.

We'll give bash a quick run to show how this works

% firejail bash

By default the bash firejail profile doesn't allow network access. Time to test it

% ping www.google.com

No success

FireJail

Let's disable the default profile so we can hit the network

% firejail –noprofile bash

Now we should be able to ping

% ping www.google.com

FireJail

Let's take a look at the default profiles that are available on the system

% ls /etc/firejail

Ok. There are a lot of them - 1,230 of them on my Debian box

You can use a profile with a program, or you can create your own either by hand or by using firejail-ui

Let's take a look at firejail-ui

FireJail-UI

% firejail-ui

So now we can invoke a gui and generate an example profile.

I still haven't played with this too much yet, but it trys to make it easy to create the profiles.

Keep in mind I suggest you use a pre-built profile as a baseline if one is available for you.

Firefox Containers

So far we've been talking about containers in the form of Docker/Podman/LXC/LXD type.

That isn't the only type of containers available.

Another example is containers that are built into Firefox.

These segment off cookies/databases/cache into separate environments

Firefox Containers/Multi User Account

This is a pretty useful plugin. It allows you to easily create separate tab types for different tasks - banking can be kept separate from gaming, shopping etc.

I personally use this with the following containers at home

- NEbraskaCERT
- Mailchimp
- Gsuite
- Regular User
- And a bunch of others

Firefox Containers/Multi User Account

Highly recommend this plug-in

Of course you are trusting that the separation between the accounts is good enough.

The old way I used to do this was by using 5 different browsers: Google Chrome/Chromium/Brave/Opera and Vivaldi, so just being able to use one browser is quite a benefit

FireFox Containers.

There are also Firefox containers available for the following uses

- Facebook - keeps cookies from Facebook separate from your other windows making it harder to track you
- Google - keeps google cookies from tracking you when you are on non-google sites. There are variants of this that keep all google together except youtube and so on
- Others exist as well Twitter and so on. There is probably a container for whatever you hit

Google also has some container extensions as well, but not as many and this will be changing with Google's upcoming plugin revisions

Sandboxie/Sandboxie Plus

Sandboxie is a simple solution for running potentially untrusted applications on Microsoft Windows.

Sandboxie Plus is the Open Source continuation of the original commercial Sandboxie program.

You can do things like restrict networking for an application, access to the filesystem and so on.

If you're using Windows this a tool you should really at least give it a try.

Containers you're not probably aware of

A lot of things are being put in containers without you doing anything.

If you're installing using applications from any software the Mac Store/ios App/Google Play Store/Windows Store/AppImage/Flatpak/Ubuntu Snaps and others these all use container type technologies over time

- Restrict access to the underlying file system
- Restrict permissions (location, networking, access to certain folders such as Photos and so on)

Containers you're not probably aware of.

OpenBSD has worked at splitting their applications into multiple parts.

As part of this the OpenSSH server has been split into multiple parts

- A privileged part that does the needed capabilities of accessing low ports and so on
- A non-privileged part that does the rest of the work

This technique has been used for a lot of their programs.

Summary

- We've come a loonnngggg way from only having chroot'd jails which we were using 40 years ago.
- The idea of least privilege is really impressive
- The tools are better. 20 years ago you could create systrace profiles for all the programs on your system and lock things down, but one change and it stopped working

Q & A

Any Questions?

Thanks for listening.

Links

Toolbox

https://docs.fedoraproject.org/en-US/fedora-silverblue/toolbox/

https://ryan.himmelwright.net/post/intro-to-toolbox/

KASM Workspaces

https://www.kasmweb.com/

Links

Distrobox

https://github.com/89luca89/distrobox/blob/main/docs/compatibility.md#supported-container-managers

Firejail

https://firejail.wordpress.com/

Links

Firefox Containers

Multi Account Containers

https://addons.mozilla.org/en-US/firefox/addon/multi-account-containers/

Facebook Containers

https://support.mozilla.org/en-US/kb/facebook-container-prevent-facebook-tracking

Links

Firefox containers

Search for available containers for Firefox

https://addons.mozilla.org/en-US/firefox/search/?q=containers