illuria SECURITY

Building a cloud in 360 seconds with FreeBSD

Antranig Vartanian

`whoami'

illuria SECURITY

- Co-founder && CEO @ illuria Security, Inc.
- HEAD @ EVNCERT
- Daemon @ Armenia BSD User Group
- Past: CTO, Systems Engineer
- FreeBSD developer and advocate
- Runs Jabber.am :-)

Disclaimer: I have no idea wtf I am talking about!

Agenda

illuria SECURITY

- FreeBSD: The Operating System
- ZFS: The File System
- VNET: The Virtual Network Stack
- Jail: Containers before it was cool
- Automation: Because we're lazy
- Demo: What could possibly go wrong

FreeBSD: Unix made with ❤️

illuria SECURITY

- Unix Operating System
- Rock-Solid. Battle-Tested. Community-Driven
- Self-Hosted / Complete
- Runs on 8 architectures and all clouds
- Oldest Democratic Running Open Source Project
- Used by multiple corporation: Apple, Netflix, Sony, NetApp, WhatsApp, Mellanox, Nginx, Microsoft, illuria Security
- By Unix people for Unix people

ZFS: The File System

illuria SECURITY

- Developed by Sun Microsystems for the Solaris operating system
- Integrity checking (data & meta-data)
- Self-healing features
- Redundancy with mirroring, RAID-Z1/2/3
- High storage capacities — up to 256 trillion yobibytes (2^128 bytes)
- Efficient storage with snapshots and copy-on-write clones
- Hardware-accelerated native encryption
- Efficient local or remote replication (zfs send, zfs recv)
- Runs on illumos, FreeBSD, NetBSD, Linux, macOS, Windows

VNET(9):

illuria SECURITY

- Network subsystem virtualization infrastructure
- Each (virtual) network stack is attached to a prison ***(Jail)***

  with vnet0 being the un-restricted default network stack of the base system.

`Jail(2)`: Containers before it was cool 😎

illuria SECURITY

- OS-Level Virtualization
- `chroot(2)` on steroids
- Each process is assigned to a Jail (pr_id)
- Each Jail has a
  - ID
  - Path
  - VNET Stack (if you want, recommended)
  - {Host}Name
  - Other Options

`Jail(2)`: Simple because we can

illuria SECURITY

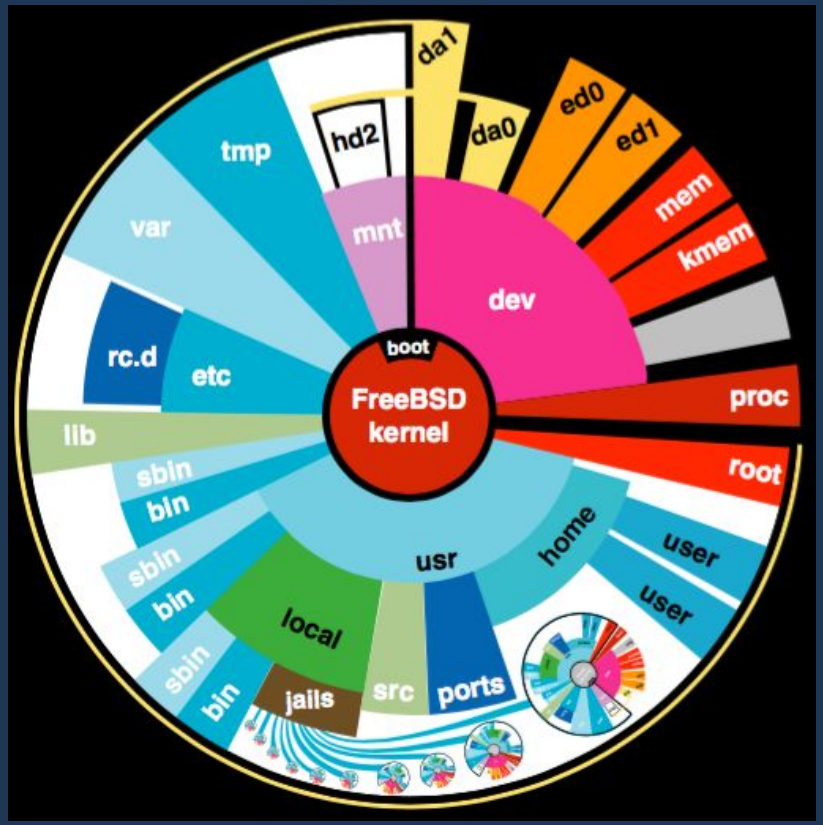

Cloud: The Harsh Reality

illuria SECURITY

- Compute
- Network
- Storage

Cloud: The Simplest Way

illuria SECURITY

- Compute → **Jail**
- Network → **VNET**
- Storage → **ZFS**
- Host → **FreeBSD**
- Automation → `/bin/sh`

demo | The base system

illuria SECURITY

```
# zfs create -o mountpoint=/usr/local/jails zroot/jails

# mkdir /usr/local/jails/.distfiles

# fetch -o \
/usr/local/jails/.distfiles/FreeBSD-12.1-RELEASE-base.txz \
https://download.freebsd.org/ftp/releases/amd64/amd64/12.2-RELEASE/base.txz

# zfs create zroot/jails/wwww

# tar xf \
/usr/local/jails/.distfiles/FreeBSD-12.1-RELEASE-base.txz\
-C /usr/local/jails/www/
```

demo | The base network

illuria SECURITY

```
# sysrc cloned_interfaces="bridge0"

# sysrc ifconfig_bridge0="inet 10.0.0.1 netmask 0xffffff00 descr jails-bridge"

# service netif start bridge0

# echo 'nat on vtnet0 inet from 10.0.0.0/24 to any -> vtnet0:0' \
  >  /etc/pf.conf && service pf start \
  && sysctl net.inet.ip.forwarding=1
```

demo | The Jail

illuria SECURITY

```
# jail -c \
    name=www\
    path=/usr/local/jails/www\
    vnet\
    host.hostname=www.illuriasecurity.com\
    persist;
```

demi | The Jail Network

illuria SECURITY

```
(host) # ifconfig epair create

(host) # ifconfig bridge0 addm epair0a

(host) # ifconfig epair0b vnet www

(host) # jexec -l www

(guest)# ifconfig epair0b inet 10.0.0.80/24 up
```

demo | Better Jail management

illuria SECURITY

```
# vi /etc/jail.conf # Next Slide for example

# sysrc jail_enable=YES

# sysrc jail_list=www

# sysrc jail_list+=oragir
```

demo | jail.conf

illuria SECURITY

```
oragir {
        $id             = "30";
        $addr           = "192.168.10.${id}";
        $mask           = "255.255.255.0";
        $gw             = "192.168.10.1";
        vnet;
        vnet.interface  = "epair${id}b";

        exec.prestart   = "ifconfig epair${id} create up";
        exec.prestart   += "ifconfig epair${id}a up descr vnet-${name}";
        exec.prestart   += "ifconfig bridge10 addm epair${id}a up";

        exec.start      = "/sbin/ifconfig lo0 127.0.0.1 up";
        exec.start      += "/sbin/ifconfig epair${id}b ${addr} netmask ${mask} up";
        exec.start      += "/sbin/route add default ${gw}";
        exec.start      += "/bin/sh /etc/rc";

        exec.poststop   = "ifconfig bridge10 deletem epair${id}a";
        exec.poststop  += "ifconfig epair${id}a destroy";

        host.hostname = "${name}.pingvinashen.am";
        path = "/usr/local/jails/${name}";
        exec.consolelog = "/var/log/jail-${name}.log";
        persist;
}
```

demo | Integration is nice

illuria SECURITY

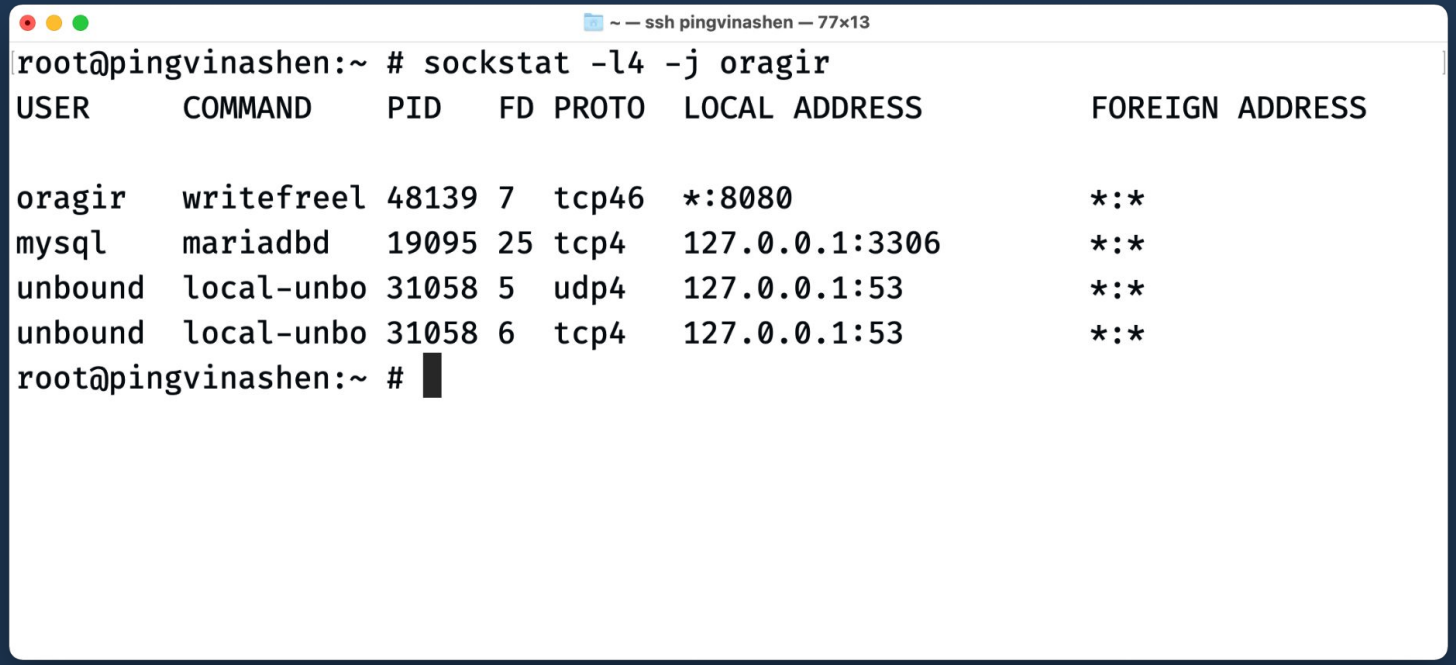

demo | Everything is Integrated

illuria SECURITY

```
(host) # pkg -j www install nginx

(host) # sysrc -j www nginx_enable=YES

(host) # ps -aux -J www

(host) # jexec www ifconfig
```

demo | All is good, save the goods

illuria SECURITY

```
# zfs snap zroot/jails/www@all_good

# zfs snap zroot/jails/www@`date -I`

(guest)# echo 'oops I did something bad'

# zfs rollback zroot/jails/www@all_good && echo 'all good!'
```

demo | Mass Deploy for multiple envs

illuria SECURITY

```
# zfs snap zroot/jails/app@base

# zfs clone zroot/jails/app@base zroot/jails/test0

# $EDITOR /etc/jail.conf
```

demo | Send/Recv

illuria SECURITY

```
(huginn)# zfs snap zroot/jails/www@good

(huginn)# zfs send zroot/jails/www@good \
  | ssh muninn zfs recv zroot/jails/www_from_huginn
```

FreeBSD Jails @ $HOME

illuria SECURITY

```
root@pingvinashen:~ # zfs list -r zroot/jails
NAME                      USED  AVAIL  REFER  MOUNTPOINT
zroot/jails              6.08G   390G   301M  /usr/local/jails
zroot/jails/bsd           438M   390G   438M  /usr/local/jails/bsd
zroot/jails/freshrss     1.25G   390G   804M  /usr/local/jails/freshrss
zroot/jails/git           965M   390G   965M  /usr/local/jails/git
zroot/jails/matterbridge  873M   390G   489M  /usr/local/jails/matterbridge
zroot/jails/psql          529M   390G   529M  /usr/local/jails/psql
zroot/jails/rss           478M   390G   478M  /usr/local/jails/rss
zroot/jails/test          361M   390G   361M  /usr/local/jails/test
zroot/jails/znc           994M   390G   994M  /usr/local/jails/znc
root@pingvinashen:~ # ifconfig -l
em0 lo0 bridge0 bridge10 bridge20 wg0 epair3a epair11a epair52a epair5a epair51a epair510a
root@pingvinashen:~ # jls
   JID  IP Address      Hostname                      Path
     1                  psql.pingvinashen.am          /usr/local/jails/psql
     2                  matterbridge.pingvinashen.am  /usr/local/jails/matterbridge
     3                  znc.bsd.am                    /usr/local/jails/znc
     4                  rss.bsd.am                    /usr/local/jails/rss
     5                  git.bsd.am                    /usr/local/jails/git
root@pingvinashen:~ # ls /usr/local/jails/matterbridge/
.cshrc     boot/      etc/       libexec/   net/       root/      tmp/
.profile   COPYRIGHT  home@      media/     proc/      sbin/      usr/
bin/       dev/       lib/       mnt/       rescue/    sys@       var/
root@pingvinashen:~ # jexec matterbridge
root@matterbridge:/ # :)
```

Patch: Better Automation is Good

illuria SECURITY

Add support for jail.d - D24570 - https://reviews.freebsd.org/D24570

```
Using /etc/jail.{jailname}.conf is nice, however it makes
/etc/ very messy if you have many jails, this patch will
help to have jail configurations in /etc/jail.conf.d

# my_jail_gen.sh >> /etc/jail.conf.d/foobar.conf
```

jailio | Because we had a bet

illuria SECURITY

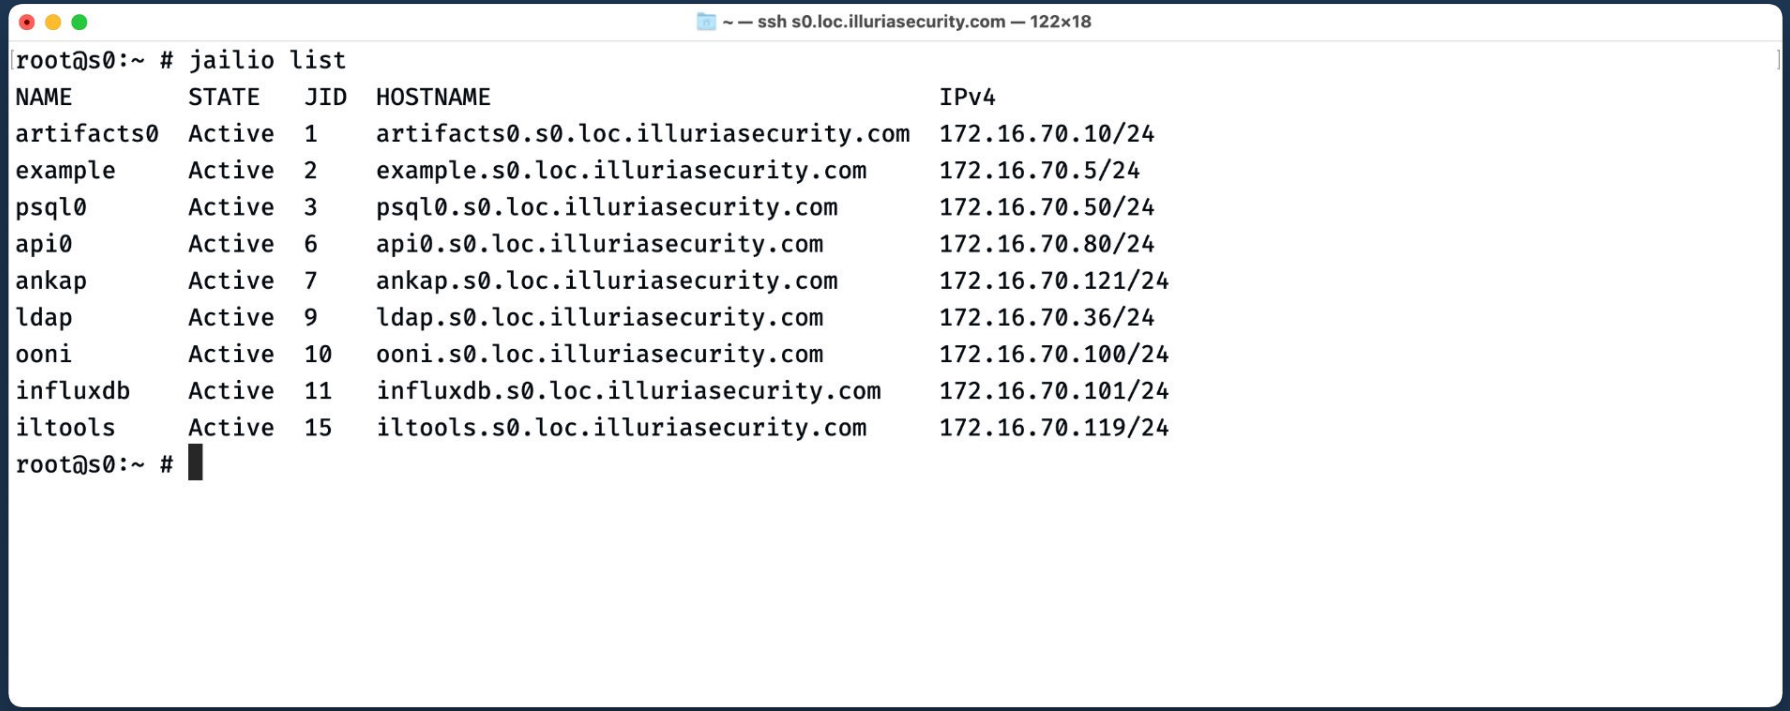

Conclusion

illuria SECURITY

- Single Operating System
- All Integrated
- ZFS saves the day
- Jail is simple
- VNET is… VNET
- Easy to automate
- Easy to maintain
- Hard to ignore

illuria SECURITY

That's all folks!

Thanks

Q&A

a@illuriasecurity.com
https://antranigv.am/