

D. Delmar Davis

None

None

None

Table of contents

1. Digithink.com.	7
2. Buildnotes	8
2.1 Building Circuit Python.	8
2.2 FreeBSD on lxd	10
2.3 Openwrt build notes (er-lite3 v19.07.7)	13
2.4 Nginx Server Build	15
2.5 LXD snapshot host	19
2.6 RockyLinuxLXDHost	20
2.7 Ansible	23
2.8 Docker vm configuration	30
2.9 Edge server configuration	38
2.10 Gitea configuration	44
2.11 Mkdocs server configuration	47
2.12 Redmine configuration	49
3. Resume	52
3.1 D Delmar Davis	52
4. Rethinkeverything	57
4.1 There is no bullet list like MY Bullet list	57
4.2 Costello, an Ubuntu 22.04 lxd 5 home server.	59
4.3 otto (OTTO) a ubuntu laptop/server.	62
4.4 pure config is out of scope of this note	63
4.5 setting up pure -> freebsd	63
4.6 /etc/rc.conf.local	63
4.7 grab rc.d file from zfs-backup2:/etc/rc.d/zpool_iscsi	63
4.8 so system will try to import zpool after iscsi has settled	63
4.9 /etc/iscsi.conf	63
4.10 view iscsi luns	63
4.11 remove luns	63
4.12 add luns	63
4.13 freebsd initiator doesn't handle multipath.	64
4.14 The geom_multipath kernel module does	64
4.15 create multipath device	64
4.16 make it survive a reboot	64
4.17	64
4.18 now you can create a zpool using the mp0 device	64

4.19	Utah	65
4.20	LXD5	66
4.21	Sense	67
5.	Serverdocs	71
5.1	Systems Documentation	71
5.2	Portland	72
5.3	Annie	73
5.4	Ansible Scripts	74
5.5	BS2020 (RE)Install	75
5.6	Bleeding Edge (old)	79
5.7	CloudServerConfiguration	80
5.8	CloudServerDocs	85
5.9	COBOL / Postgres / Open Enterprise/Government (Rough Draft)	86
5.10	Containership Creation	87
5.11	Updating Hosts Notes	88
5.12	Old DL380 Raid Notes	91
5.13	DeeDee	101
5.14	BS2020 LXC to LXD Notes	105
5.15	DDRescue Notes	107
5.16	Docker Installation on FranklinOnce the LXD container for docker was built out I followed the directions at docker.com to install docker-ce	108
5.17	DL380 Raid Notes	109
5.18	Esp8266	117
5.19	Feurig	118
5.20	FocalNotes	119
5.21	Gold Coast	120
5.22	Goodbye Openstack	123
5.23	from https://docs.openstack.org/devstack/latest/guides/lxc.html	125
5.24	Permit access to /dev/loop*	125
5.25	Setup access to /dev/net/tun and /dev/kvm	125
5.26	Networking	125
5.27	IPV4_ADDRS_SAFE_TO_USE=172.31.1.0/24	126
5.28	Hardening LEDE	127
5.29	Overview	129
5.30	Joey Snippet	131
5.31	LXDContainerWithDockerNotes	132
5.32	FIRST IMPRESSIONS:	133
5.33	This file describes the network interfaces available on your system	133

5.34	and how to activate them. For more information, see interfaces(5).	133
5.35	The loopback network interface	133
5.36	Source interfaces	134
5.37	Please check /etc/network/interfaces.d before changing this file	134
5.38	as interfaces may have been defined in /etc/network/interfaces.d	134
5.39	See LP: #1262951	134
5.40	change this after first instantiation	135
5.41	127.0.0.1 localhost	135
5.42	The following lines are desirable for IPv6 capable hosts	135
5.43	Imagebuilder notes	136
5.44	OPENVPN on LEDE Notes	138
5.45	OpenWRT Notes	142
5.46	LEDE EA3500 Note	144
5.47	LEDE build (old)	145
5.48	LEDE E900 Notes	146
5.49	LEDE Remote Syslog	147
5.50	Lets Encrypt Certificates	149
5.51	MigrateUsers	151
5.52	Migrating Services to LXD	152
5.53	default server and configuration	154
5.54		154
5.55	virtualhosts	154
5.56		154
5.57	disable php	154
5.58	Mullein	158
5.59	New Trac Container	159
5.60	Nigel	165
5.61	NotesAddingAnsibleToContainerCreation	166
5.62	Notes: Automating Container Updates	167
5.63	Buster Notes	169
5.64	HP Z400 notes	171
5.65	NotesOnAppleTalk3vsUbuntu	174
5.66	DL380 Raid Bios notes	175
5.67	ILO3 Notes	177
5.68	Irac6 Notes	180
5.69	LXD FIRST IMPRESSIONS:	182
5.70	NotesOnUbuntu18.04	185
5.71	HOLY FUCKING AWESOME!!!!	188

5.72	OpenVPN on LEDE (Fail)	192
5.73	https://lede-project.org/docs/user-guide/openvpn.server	193
5.74	OpenWRT Notes	195
5.75	LEDE on EA3500	197
5.76	Building old LEDE firmware	198
5.77	OpenWRT E900 Firmware Build	199
5.78	Server Modernization	200
5.79	Tasks: Accessing Hosts	205
5.80	links.... (tbd)	207
5.81	PlatformIO	208
5.82	RecentChanges	209
5.83	Redmine Install	210
5.84	Foobarred zfs filesystem	211
5.85	Start Using the F words.	213
5.86	SuspectDevices	214
5.87	System Updates (for gihon)	215
5.88	TaskAddGitHubRepo	218
5.89	TaskAddLxdContainerWithAnsible	220
5.90	TaskCreatingNewContainers	222
5.91	Task: Dual Proxy Configuration	227
5.92	Fast Forward	228
5.93	Install Ansible	232
5.94	Squid Caching Server	236
5.95	Task: Split ZF Mirror	238
5.96	Task: ZFS Disk Replacement	242
5.97	zfs list and check for larger disk pool	244
5.98	zpool set autoexpand=on devel	245
5.99	zpool online -e devel scsi-xxxxxxxxxxxxxxxxxxxxxx	245
5.100	zpool online -e devel scsi-yyyyyyyyyyyyyyyyyyyy	245
5.101	zpool set autoexpand=off devel	245
5.102	Ubuntu18.04Notes	246
5.103	Ubuntu LTS Email Server Setup	249
5.104	VideoRanchCloudServerConfiguration	254
5.105	fixing	259
5.106	ZFS Disk Replacement	261
5.107	HOLY FUCKING AWESOME!!!!	264
5.108	ZFS Mirrored data on existing file server	268
5.109	ZFS IS ALL THE RAGE	270

5.110 kb2018 install bash history.	272
5.111 Unifi install	277

1. Digithink.com.

I'll be 72 when unix time ends.



After cutting my teeth on Dec VMS systems in college, I worked on my first unix systems in the late 80s. I am still working in linux. Collected here are some notes past and present.

And here is my [resume](#)

Did you think you'd be doing this for more than 30 years?

Maybe its time to [rethink everything](#)

2. Buildnotes

2.1 Building Circuit Python.

I was looking at adafruit's circuit python as a possible platform to replace the missing link using either the rpi2040 or the esp32s2. Unfortunately the usb stacks supported are different so I knew I would have to build what I want from scratch. I am not sure that this isn't a rathole.

<https://learn.adafruit.com/building-circuitpython/build-circuitpython>

Building on Sandbox (ubuntu focal)

Circuitpython for most processors requires more than a few dependencies.

```
sudo apt-get install build-essential git gettext uncrustify python3-pip
sudo apt-get install python3-setuptools cmake ninja-build ccache libffi-dev libssl-dev dfu-util libusb-1.0-0
sudo pip3 install cascadetoml
sudo apt-get install gcc-arm-none-eabi
```

This should allow you to check out the source code for circuitpython as a user. As per their suggestions I forked the repo first.

```
git clone git@github.com:suspect-devices/circuitpython.git
cd circuitpython/
git submodule sync --quiet --recursive
git submodule update --init
```

Once there you can install further dependencies.

```
sudo pip3 install -r requirements-dev.txt
sudo make -C mpy-cross
```

At which point you can build firmware for most targets.

```
cd ports/raspberrypi/
make BOARD=adafruit_feather_rp2040
cd build-adafruit_feather_rp2040/
```

ADDING THE ESPRESSIF TOOLCHAIN AND IDF

```
cd /home/feurig/circuitpython/ports/esp32s2#
sudo esp-idf/install.sh
```

First attempt.

On my first attempt I was only interested in building for the esp32s2. I ran everything as root including the builds. It is better to separate privileges with anything this large where you are building for a different target.

```
root@viva:# apt-get install git wget flex bison gperf python3 python3-pip python3-setuptools cmake ninja-build ccache libffi-dev libssl-dev dfu-util libusb-1.0-0
...
root@viva:/home/feurig# mkdir -p ~/esp
root@viva:/home/feurig# cd ~/esp
root@viva:~/esp# git clone --recursive https://github.com/espressif/esp-idf.git
Cloning into 'esp-idf'...
...
root@viva:~/esp# cd esp-idf/
root@viva:~/esp/esp-idf# ./install.sh
...
All done! You can now run:

  . ./export.sh

root@viva:~/esp/esp-idf# . ./export.sh
...
root@viva:~# cd /home/feurig/circuitpython/
root@viva:/home/feurig/circuitpython# cd ports/
root@viva:/home/feurig/circuitpython/ports# cd esp32s2/
root@viva:/home/feurig/circuitpython/ports/esp32s2# make BOARD=unexpectedmaker_feathers2
...
Wrote 2601984 bytes to build-unexpectedmaker_feathers2/firmware.uf2
```



```
root@viva:/home/feurig/circuitpython/ports/esp32s2#
```

Then as a comparison I built for the adafruit feather rp2040.

```
root@viva:/home/feurig/circuitpython/ports# cd raspberrypi/
root@viva:/home/feurig/circuitpython/ports/raspberrypi# ls boards/
adafruit_feather_rp2040  pimoroni_keybow2040  raspberry_pi_pico
adafruit_itsybitsy_rp2040  pimoroni_picosystem  sparkfun_pro_micro_rp2040
adafruit_qtpy_rp2040      pimoroni_tiny2040    sparkfun_thing_plus_rp2040
root@viva:/home/feurig/circuitpython/ports/raspberrypi# make BOARD=adafruit_feather_rp2040
Use make V=1, make V=2 or set BUILD_VERBOSE similarly in your environment to increase build verbosity.
QSTR updated
Traceback (most recent call last):
  File "gen_stage2.py", line 2, in <module>
    import cascadetoml
ModuleNotFoundError: No module named 'cascadetoml'
root@viva:/home/feurig/circuitpython/ports/raspberrypi# pip3 install cascadetoml
...
root@viva:/home/feurig/circuitpython/ports/raspberrypi# make BOARD=adafruit_feather_rp2040
...
Wrote 1259520 bytes to build-adafruit_feather_rp2040/firmware.uf2
root@viva:/home/feurig/circuitpython/ports/raspberrypi#
```

2.2 FreeBSD on lxd

LXD 4.0 allows for the creation of VM's based on qemu. This allows us to create "virtual machines" capable of running non linux operating systems such as FreeBSD (or god forbid WindBlows). So let's look at adding a freebsd 12.3 box to our setup.

2.2.1 Create an empty vm.

Based on the examples I was able to find we start by creating an empty vm and then tweak on a few of the parameters (raw.apparmor and raw.qemu). While there i adjust the nic (I am sure that all of this could be done on the init line). After that it's pretty straight forward.

```

root@bs2020:/home/feurig# lxc --empty --vm -c limits.cpu=4 -c limits.memory=4GB -c security.secureboot=false -n br0
Creating henry
root@bs2020:/home/feurig# lxc config device add henry install disk source=/home/feurig/FreeBSD-12.2-RELEASE-amd64-dvd1.iso
Device install added to henry
root@bs2020:/home/feurig# lxc config edit henry
architecture: x86_64
config:
  limits.cpu: "4"
  limits.memory: 4GB
  security.secureboot: "false"
  ## tweak apparmor/qemu settings
  raw.apparmor: /home/feurig/** rwx,
  raw.qemu: -boot menu=on -machine pc-q35-2.6
  volatile.apply_template: create
  volatile.br0.hwaddr: 00:16:3e:ab:07:4e
  volatile.eth0.hwaddr: 00:16:3e:87:3c:b1
devices:
  eth0:
    nictype: bridged
    parent: br0
    type: nic
ephemeral: false
profiles:
- default
stateful: false
description: "FreeBSD 12.3 test box"
root@bs2020:/home/feurig# lxc start henry --console

```

[illegible]

I found that, on at least one of my servers, the console would not come up with the dual "Cons:" setting. Serial worked just fine.

2.2.2 Next Steps (sudo, ssh, hardening, usw)

In order to have the server play well with our environment I install the following packages using `pkg` (`sudo`, `nano`, `bash`, `bash-completion`, `python37`) as well as manually adding admin users. At some point it would be nice to use cloud-init or if that is unworkable ansible for the initial configuration.

```
[root@henry /usr/home/feurig]# pkg info
bash-5.1.4_1      GNU Project's Bourne Again Shell
bash-completion-2.11,2  Programmable completion library for Bash
gettext-runtime-0.21  GNU gettext runtime libraries and programs
indexinfo-0.3.1    Utility to regenerate the GNU info page index
libffi-3.3.1       Foreign Function Interface
nano-5.5            Nano's ANOTHER editor, an enhanced free Pico clone
pkg-1.16.3         Package manager
```

py37-pip-20.2.3	Tool for installing and managing Python packages
py37-setuptools-44.0.0	Python packages installer
python37-3.7.10	Interpreted object-oriented programming language
readline-8.1.0	Library for editing command lines as they are typed
sudo-1.9.6p1	Allow others to run commands as root

Nano and bash are a personal preference of mine.

```
feurig@henry:~$ sudo bash
Password:
[root@henry /usr/home/feurig]# chpass -s /usr/local/bin/bash feurig
[root@henry /usr/home/feurig]# chpass -s /usr/local/bin/bash joe
```

2.2.3 Setting up Ansible on BSD

In addition to installing python ssh and an admin user needs to be set up as lxd does not "lxc exec" directly to virtual machines.

```
[root@henry /usr/home/feurig]# visudo
... comment out this
# root ALL=(ALL) ALL
... and uncomment out this
%wheel ALL=(ALL) ALL
...
[root@henry /usr/home/feurig]# adduser ansible
... add ansible to wheel group ...
[root@henry /usr/home/feurig]# su - ansible
ansible@henry:~$ ssh-keygen
ansible@henry:~$ cat >> .ssh/authorized_keys
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCsxhi6P1Ssin8QjEMlm+9W1L5ncRqejnw78z/
yhQLwCU2av3+vAzPFDKi7CTm2iqeRoNYsKx4IaNYK9t+zQ00sEXjzIz5+uCNODbNaw4pMtaHcwsaYDCdG90iXuFa7qWndAvSJjXQR6t1pygdw/tbbsGN0//zq71j9ChitXJQR0YYCYwa4MaB6Srn/
Zpkhfut10P56XMo15F+0YD+oS/IqJp/QTH6Q9LzVh+HKI9rdhEqEsrNZsaQw6UZ8JrFRYmJWzcFlqztv2qBv/BdStWbJGMBDTDN0S9f9wKts43lkZGYgSyZo80NLmq4oXJanuN00w0BeRtMyX+HUEmgh root@kb2018
<ctrl-D>
```

Adding the become password to the ansible servers vault is described [here](#)

2.2.4 Adding an update.sh script.

The field expedient way

For quick and dirty's sake we add the following to /usr/local/bin/update.sh which could easily be added to the generalized shell since we have decided that bash is ok. It also might be ok to check if a reboot is necessary.

```
[root@henry /usr/home/feurig]# cat /usr/local/bin/update.sh
#!/bin/sh
freebsd-update fetch install
pkg upgrade
```

And then life is good and our new pets are equally loved. (including the centos 7 result for sag)

```
root@kb2018:/etc/ansible/python# ansible pets -m raw -a "update.sh"
shelly | CHANGED | rc=0 >>
...
henry | CHANGED | rc=0 >>

src component not installed, skipped
Looking up update.FreeBSD.org mirrors... 2 mirrors found.
Fetching metadata signature for 12.2-RELEASE from update1.freebsd.org... done.
Fetching metadata index... done.
Inspecting system... done.
Preparing to download files... done.

No updates needed to update system to 12.2-RELEASE-p8.
No updates are available to install.
Updating FreeBSD repository catalogue...
FreeBSD repository is up to date.
All repositories are up to date.
Checking for upgrades (1 candidates): 100%
Processing candidates (1 candidates): 100%
Checking integrity... done (0 conflicting)
Your packages are up to date.
Shared connection to henry closed.

keynes | CHANGED | rc=0 >>
----- begin updating keynes -----
yum upgrade.
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
* base: mirrors.cat.pdx.edu
* extras: mirror.web-ster.com
```

```
* updates: mirror.keystealth.org
No packages marked for update
===== done =====
Failed to set locale, defaulting to C
...
naomi | CHANGED | rc=0 >>
----- begin updating naomi -----
Get:1 http://security.ubuntu.com/ubuntu bionic-security InRelease [88.7 kB]
Hit:2 http://archive.ubuntu.com/ubuntu bionic InRelease
```

Making it right.

Once we make peace with installing/enforcing bash on a freebsd box then we can add the freebsd update to our multi platform [update.sh](#) (again someday deployed by lxd/cloud-init).

```
[root@henry /usr/home/feurig]# ln -s /usr/local/bin/bash /bin/
[root@henry /usr/home/feurig]# nano /usr/local/bin/update.sh
#!/bin/bash
# update.sh for debian/ubuntu/centos/suse/freebsd/pihole
# https://bitbucket.org/suspectdevicesadmin/ansible/src/master/files/update.sh
# (copyleft 2021) don@suspectdevices.com
echo ----- begin updating `uname -n` -----
if [ -x "$(command -v apt-get)" ]; then
    echo Updating system
    apt-get update
    apt-get -y dist-upgrade
    apt-get -y autoremove
fi
if [ -x "$(command -v yum)" ]; then
    echo yum upgrade.
    yum -y upgrade
fi
if [ -x "$(command -v pihole)" ]; then
    echo Updating pihole.
    pihole -up
fi
if [ -x "$(command -v zypper)" ]; then
    echo zypper dist-upgrade.
    zypper -y dist-upgrade
fi
if [ -x "$(command -v freebsd-update)" ]; then
    echo Updating freebsd base
    freebsd-update fetch install
    echo Updating freebsd packages
    pkg upgrade
fi
echo ===== done =====
```

2.2.5 To do.

- Look at restricting ansible's ssh access to hosts on the admin lan (as is done for bs2020).
- Add virtual machines to nightly backups (currently only containers).

```
root@kb2018:/# lxc snapshot henry 2021-06-05
root@kb2018:/# lxc move henry/2021-06-05 bs2020:Spare-henry-2021-06-05
root@kb2018:/# lxc stop bs2020:Spare-henry-2021-06-05
Error: The instance is already stopped
```

Linkdump.

- <https://forum.netgate.com/topic/154906/how-to-install-pfsense-on-lxc-vm-qemu>
- <https://discuss.linuxcontainers.org/t/lxc-vm-running-freebsd-cant-see-hard-disk/8214/14>
- <https://download.freebsd.org/ftp/releases/ISO-IMAGES/12.2/FreeBSD-12.2-RELEASE-amd64-dvd1.iso>
- <https://docs.freebsd.org/en/books/handbook/bsdinstall/#bsdinstall-start>

2.3 Openwrt build notes (er-lite3 v19.07.7)

This build is based on our [build of openwrt for the home lan](#).

Since the er-lite uses a usb boot drive we are no longer constrained by the usual space restrictions. For that reason we have room to add python/ansible to managing this host. We still trim the os as much as possible in order to minimize the attack surfaces available to the wild. (ipv6 and luci aren't needed here for instance)

Objectives

- pre-build os hardening.
- python3 for ansible management
- wireguard
- ipv4 only
- git based configuration management (if possible)

Changes since 19.07.3

- Somewhere since v19.07.3 the shadow password file has been rolled into the OS. (*So adduser instead of shadow_adduser add etc...*).
- The built in logger conflicts with syslog_ng so I hope that means we can pipe our logs off to another system.

2.3.1 Pre hardening and initial configuration during build.

In our deployment the router is maintained externally. For this reason direct login to the router as root is disabled and sudo enabled accounts are installed. These accounts connect using ssh keys and escalate privileges with their passwords. The root account is locked and ssh access is allowed from the wan port. The process for this is documented [here](#) This configuration is added to the build under the files directory where they are copied into the root filesystem of the target. The box then comes up pre configured and pre-hardened. One kludge used here is to add an rc.local which changes the users home directories to be owned by them. Otherwise the ssh keys will not have the correct permissions. The files directory is maintained in a [private git repository](#). Additionally /etc/sudoers, /etc/rc.local, and /home are added to /etc/sysupgrade.conf in order to preserve them during sysupgrade.

Changes made to the os.

- Added sudo admin accounts and locked the root account.
- Locked the console. The console out of the box is wide open normally that wouldn't be a problem since the serial ports on most routers aren't exposed.

```
# nano /etc/config/system
...
option ttylogin '1'
...
```

- Added python3 and tested pip for future work
- Added and configured postfix as a satellite system
- removed symlink to /etc/resolv.conf

I have no idea why everyone has got to mess this up.

ADDITIONAL PACKAGES

- sudo
- adduser

- nano, monit
- git, git-http
- postfix

build history

```
git clone -b v19.07.7 https://github.com/openwrt/openwrt.git
cd openwrt
mv files /tmp/
git clone git@bitbucket.org:suspectdevicesadmin/goodknight-configuration.git files
cp files/lede19.07.7-erlite-ext4.diffconfig .config
make defconfig
./scripts/feeds update -a
./scripts/feeds install -a
make -j8 v=sc download world
mv bin/targets/octeon/generic/openwrt-octeon-erlite-ext4-sysupgrade.tar.gz ~/firmware/lede19.07.7-erlite-ext4.tgz
./scripts/diffconfig.sh > ~/firmware/lede19.07.7-erlite-ext4.diffconfig
./scripts/diffconfig.sh > files/lede19.07.7-erlite-ext4.diffconfig
cd files
git commit -a -m "update diffconfig in repo"
git push
```

Link Dump

- <http://www.digithink.com/serverdocs/HardeningLEDE/>
- <https://oldwiki.archive.openwrt.org/doc/howto/serial/console.password>
- <https://blog.suspectdevices.com/blahg/openwrt/lede-19-07-on-the-ubiquiti-er-lite3/>

2.4 Nginx Server Build

After lighthttpd left me with a broken configuration during the last round of updates, I started looking at nginx to serve the static sites previously served (www.digithink.com[this site], www.busholini.org). Since it and apache are both supported by the eff's certbot I was hoping that the automatic configuration and renewal features would work. And they did.

2.4.1 Nginx install.

Broke it the first try.

Nginx's debian package installs a default web server configuration which breaks if ipv6 is disabled. This breaks at the package installation. Not cool at all.

THE BROKEN.

```
Job for nginx.service failed because the control process exited with error code.
See "systemctl status nginx.service" and "journalctl -xe" for details.
invoke-rc.d: initscript nginx, action "start" failed.
● nginx.service - A high performance web server and a reverse proxy server
   Loaded: loaded (/lib/systemd/system/nginx.service; enabled; vendor preset: enabled)
   Active: failed (Result: exit-code) since Sat 2021-05-08 09:34:46 PDT; 8ms ago
     Docs: man:nginx(8)
  Process: 10025 ExecStartPre=/usr/sbin/nginx -t -q -g daemon on; master_process on; (code=exited, status=1/FAILURE)

May 08 09:34:46 guenter systemd[1]: Starting A high performance web server and a reverse proxy server...
May 08 09:34:46 guenter nginx[10025]: nginx: [emerg] socket() [::]:80 failed (97: Address family not supported by protocol)
May 08 09:34:46 guenter nginx[10025]: nginx: configuration file /etc/nginx/nginx.conf test failed
May 08 09:34:46 guenter systemd[1]: nginx.service: Control process exited, code=exited, status=1/FAILURE
May 08 09:34:46 guenter systemd[1]: nginx.service: Failed with result 'exit-code'.
May 08 09:34:46 guenter systemd[1]: Failed to start A high performance web server and a reverse proxy server.
```

THE FIX.

To fix this we correct the bad configuration file that was installed and reinstall the package. I could also have (dpkg -a --configure)d here.

```
root@guenter:/etc/nginx# nano sites-available/default
... comment out the [::]:80 ...
#listen [::]:80 default_server;
...
root@guenter:/etc/nginx# apt-get install certbot nginx
Reading package lists... Done
Building dependency tree
Reading state information... Done
nginx is already the newest version (1.18.0-0ubuntu1).
certbot is already the newest version (0.40.0-1ubuntu0.1).
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
2 not fully installed or removed.
After this operation, 0 B of additional disk space will be used.
Do you want to continue? [Y/n]
Setting up nginx-core (1.18.0-0ubuntu1) ...
Setting up nginx (1.18.0-0ubuntu1) ...
```

Adding content from other server.

Now that we have a working server we give it some content.

COPY CONTENT FROM LIGHTTPD SERVER.

```
root@guenter:/var/www# tar -xzf www.tgz
```

NAMED VIRTUAL HOSTS

```
# nano /etc/nginx/sites-available/default

server {
    listen 80;
#    listen [::]:80;
    server_name www.digithink.com;
    root /var/www/digithink/site;
    index index.html;
    location / {
        try_files $uri $uri/ =404;
```

```

    }
}

server {
    listen 80;
#    listen [::]:80;
    server_name www.busholini.org;
    root /var/www/busholini/www;
    index index.html;
    location / {
        try_files $uri $uri/ =404;
    }
}

```

Add Certificate from Lets Encrypt.

Even though this site is static and public we still want to add SSL to the site to prevent the content from being altered along the way.

CERTBOT ACTUALLY WORKED AS ADVERTIZED FOR THE FIRST TIME.

The first couple of web pages I found on the web described the automagic creation and configuration of certificate and once I replaced python-certbot-nginx with python3-cerbot-nginx things actually went brilliantly. No more --manual reinstallation.

```

root@guenter:/var/www# apt-get install python3-certbot-nginx
Reading package lists... Done
...
Setting up python3-certbot-nginx (0.40.0-0ubuntu0.1) ...
root@guenter:/var/www# certbot --nginx -d digithink.com -d www.digithink.com
Saving debug log to /var/log/letsencrypt/letsencrypt.log
Plugins selected: Authenticator nginx, Installer nginx
Enter email address (used for urgent renewal and security notices) (Enter 'c' to
cancel): don@digithink.com

-----
Please read the Terms of Service at
https://letsencrypt.org/documents/LE-SA-v1.2-November-15-2017.pdf. You must
agree in order to register with the ACME server at
https://acme-v02.api.letsencrypt.org/directory
-----
(A)gree/(C)ancel: A

-----
Would you be willing to share your email address with the Electronic Frontier
Foundation, a founding partner of the Let's Encrypt project and the non-profit
organization that develops Certbot? We'd like to send you email about our work
encrypting the web, EFF news, campaigns, and ways to support digital freedom.
-----
(Y)es/(N)o: Y
Obtaining a new certificate
Performing the following challenges:
http-01 challenge for digithink.com
http-01 challenge for www.digithink.com
Waiting for verification...
Cleaning up challenges
Deploying Certificate to VirtualHost /etc/nginx/sites-enabled/default
Deploying Certificate to VirtualHost /etc/nginx/sites-enabled/default

Please choose whether or not to redirect HTTP traffic to HTTPS, removing HTTP access.
-----
1: No redirect - Make no further changes to the webserver configuration.
2: Redirect - Make all requests redirect to secure HTTPS access. Choose this for
new sites, or if you're confident your site works on HTTPS. You can undo this
change by editing your web server's configuration.
-----
Select the appropriate number [1-2] then [enter] (press 'c' to cancel): 2
Redirecting all traffic on port 80 to ssl in /etc/nginx/sites-enabled/default
Redirecting all traffic on port 80 to ssl in /etc/nginx/sites-enabled/default

-----
Congratulations! You have successfully enabled https://digithink.com and
https://www.digithink.com

You should test your configuration at:
https://www.ssllabs.com/ssltest/analyze.html?d=digithink.com
https://www.ssllabs.com/ssltest/analyze.html?d=www.digithink.com
-----

IMPORTANT NOTES:
- Congratulations! Your certificate and chain have been saved at:
  /etc/letsencrypt/live/digithink.com/fullchain.pem
  Your key file has been saved at:
  /etc/letsencrypt/live/digithink.com/privkey.pem
  Your cert will expire on 2021-08-06. To obtain a new or tweaked
  version of this certificate in the future, simply run certbot again
  with the "certonly" option. To non-interactively renew *all* of
  your certificates, run "certbot renew"

```



```
- Your account credentials have been saved in your Certbot
configuration directory at /etc/letsencrypt. You should make a
secure backup of this folder now. This configuration directory will
also contain certificates and private keys obtained by Certbot so
making regular backups of this folder is ideal.
- If you like Certbot, please consider supporting our work by:

    Donating to ISRG / Let's Encrypt:  https://letsencrypt.org/donate
    Donating to EFF:                  https://eff.org/donate-le
```

```
root@guenter:/var/www#
```

Finishing up.

Part of this servers job is to serve this documentation. The static content is generated from markdown maintained in a git repository. To get the html we use mkdocs as described in [this document](#).

```
Hoffa-6:Documents don$ cd /Volumes/TheFlatField/static/digithink/docs/
Hoffa-6:docs don$ git add buildnotes/nginx-server-build.md
Hoffa-6:docs don$ git commit -a -m"add nginx server docs"
[main 44590d5] add nginx server docs
 1 file changed, 188 insertions(+)
 create mode 100644 docs/buildnotes/nginx-server-build.md
Hoffa-6:docs don$ git push
...
To github.com:feurig/digithink.git
 8d8f2c1..44590d5  main -> main
Hoffa-6:docs don$
```

In order to pull the content we need to add the ssh key to the github repository.

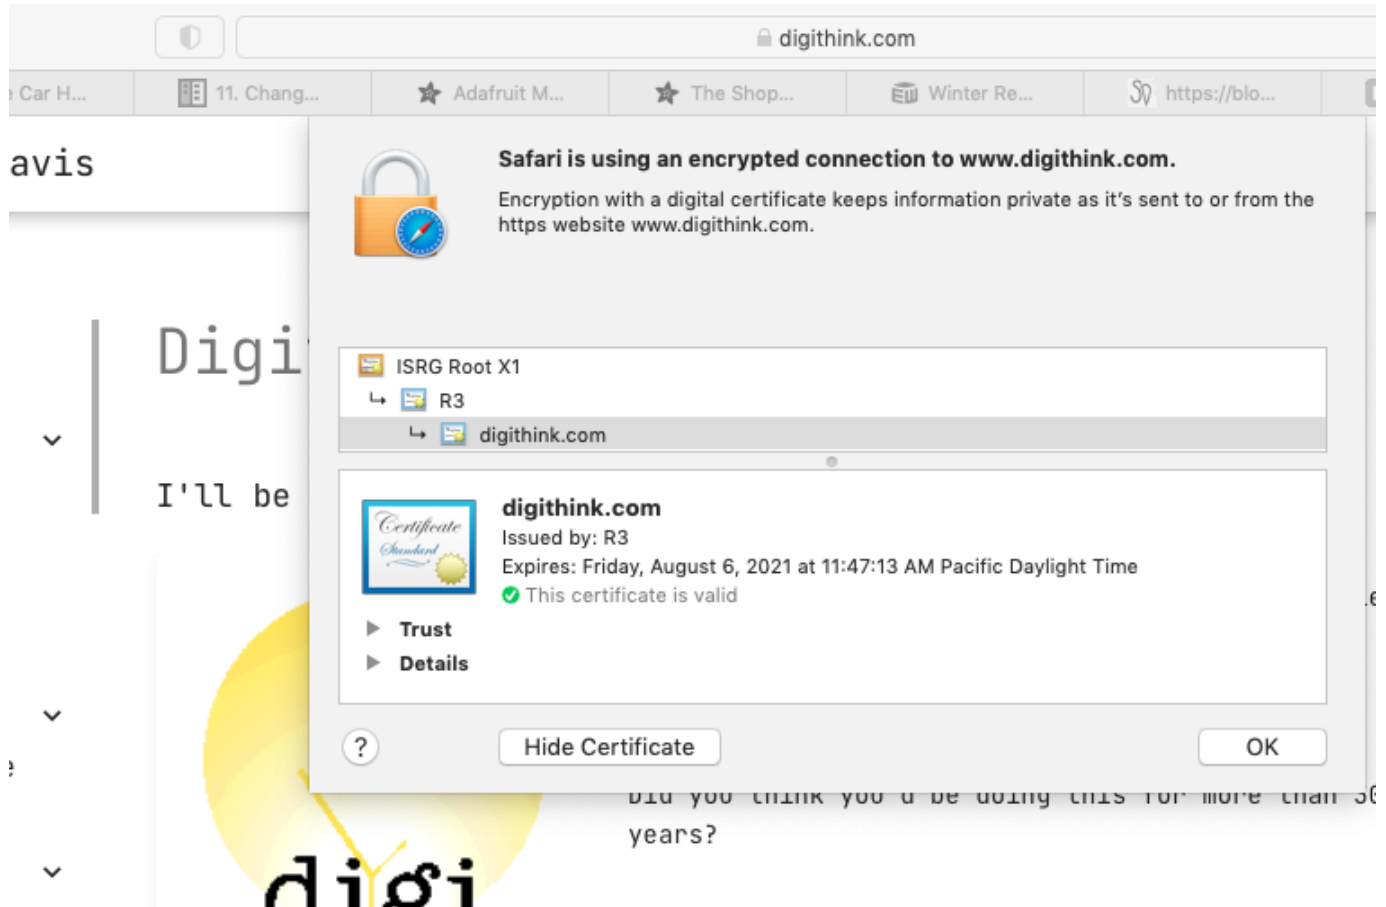
```
root@guenter:/var/www/digithink/docs# git pull
The authenticity of host 'github.com (192.30.255.112)' can't be established.
RSA key fingerprint is SHA256:nThbg6kXUpJWGl7E1IGOCspRomTxdCARLviKw6E5SY8.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added 'github.com,192.30.255.112' (RSA) to the list of known hosts.
git@github.com: Permission denied (publickey).
fatal: Could not read from remote repository.

...
root@guenter:/var/www/digithink/docs# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
...
root@guenter:/var/www/digithink/docs# cat /root/.ssh/id_rsa.pub
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQGDjUUSJi9A0kqLW85EAUkmRGGBdRiACSKhtOTNiF4Twf/PQM/ViGIAj/nkA97fdvVRpA93f7C/9hXPgd+ZCe3S5mm2KN8yxi6pqq+y6AZ/
vN58c3zeHT1YQerqzRM9WP59bQ0q2uulPh65BmcMwqlhLY+wHW5c+mIUyrW0ictg0gT8QzZdJ95hFsUkkfQWi90PID1MmiWdNh0KPKqUHTjbdv4VLjg60wggguRUJd20GRxodRON00uWqQK7A4JKqVF3LYkSym0R+xWmFDaskxmTr
+1lZGP4AigCkQJ8/3pUex7ccJoFcvhbJ8e1KGRpnlL9/BVm3baZb8iATZMb8puZoMxE/kutM8nuhP+pj004iU2QXXl62xs= root@guenter
```

Then we can regenerate the content

```
root@guenter:/var/www/digithink/docs# git pull
Warning: Permanently added the RSA host key for IP address '192.30.255.113' to the list of known hosts.
...
create mode 100644 docs/buildnotes/nginx-server-build.md
root@guenter:/var/www/digithink/docs# cd ..
root@guenter:/var/www/digithink# ls
docs  mkdocs.yml  site
root@guenter:/var/www/digithink# nano mkdocs.yml
root@guenter:/var/www/digithink# mkd
mkdir  mkdocs
root@guenter:/var/www/digithink# git submodule sync
Synchronizing submodule url for 'docs/buildnotes/ansible'
Synchronizing submodule url for 'docs/buildnotes/edge-server-configuration'
root@guenter:/var/www/digithink# mkdocs build
INFO - Cleaning site directory
INFO - Building documentation to directory: /var/www/digithink/site
INFO - Number headings up to level 3.
INFO - Generate a table of contents up to heading level 2.
INFO - Generate a cover page with "default_cover.html.j2".
INFO - Converting <img> alignment(workaround).
INFO - Rendering for PDF.
INFO - Output a PDF to "/var/www/digithink/site/pdf/document.pdf".
INFO - Converting 93 articles to PDF took 41.1s
INFO - Documentation built in 43.43 seconds
root@guenter:/var/www/digithink# chown -R www-data:www-data site/
```

And life is good.



Link Dump

- <https://webhostinggeeks.com/howto/static-website-configuration-for-nginx/>
- <https://docs.nginx.com/nginx/admin-guide/web-server/serving-static-content/>
- <https://medium.com/@jasonrigden/how-to-host-a-static-website-with-nginx-8b2dd0c5b301>
- <https://www.digithink.com/buildnotes/mkdocs-server-configuration/mkdocs-server-configuration/>
- <https://www.digitalocean.com/community/tutorials/how-to-secure-nginx-with-let-s-encrypt-on-ubuntu-20-04>

2.5 LXD snapshot host

2.5.1 Index

2.6 RockyLinuxLXDHost

2.6.1 ~~Rocky Linux 8.4~~ Ubuntu 21.04 (Hirsute Hippo) LXD host. (original AOC2024 install notes)

I had hoped that the new community supported downstream operating system based on Red Hats Enterprize Linux could take my work out of the nightmare of FC26 and Centos7. Since they run Dell Hardware, some of which is as old as our server hardware I thought it was an opportunity to test it as way forward. It failed miserably. But better to fail here where we aren't taking others with us and in ways that make us see the errors in our way.

- The original text in this document is plain text.
- ~~Where the text is no longer relevant its struck through.~~
- *Corrections are in italics*
- The new document for the server build will be at <https://www.digithink.com/buildnotes/LXD-snapshot-host/>

While bs2020 the candiate and kb2018 our Governer have served us well for the last few years its time for something new. Something community sourced, smart, and revolutionary. AOC2024

If Ubuntu's adoption of LXD and ZFS and other innovations are to mean anything they have to be separated from both debian (it's technical underpinnings) and Canonical (it's obnoxiously "freindly" commercial counterpart) or we will be captive to its "charms"(1). Meanwhile, the rpm based world has been completely paralyzed by Redhat's inability or unwillingness to provide a downstream open source project to use as a standard. This has created the disaster that is fedora 2x and the longest currently supported linux operating system ever (Centos 7 at a proposed 12 years). Redhat's choice to ditch Centos 8 and use the open source community to beta test their new features can not be described politely (2).

It ~~is~~ was our intention to support the community as it tells Redhat where to go while insuring that Debian and Ubuntu's innovations do not go to waste. Therefore, our new server ~~will run Rocky Linux~~ *would have run Rocky Linux but because bug for bug means that it won't run on 10 year old enterprise class hardware, it will run the latest Ubuntu server release* to create a robust and flexible server which will compliment the work done by our Ubuntu LTS based server.

Original goals.

- ~~Take advantage of Rocky Linux's downstream bug for bug compatibility with RHEL8~~
 - Dell's support for its hardware is limited to commercial operating systems. Attempting to get their tools (raid, idrac, configuration etc) wedged into ubuntu is like needing a root canal. *Apparently, in RHEL 8, Redhat and Dell only support what they are currently selling.*
- Use the tools that Ubuntu/Canonical has been supporting for virtualization.
 - LXD for both VMs and lxc based containers.
 - zfs, cause it rules.
- Leave our comfort zones.
 - And still do production quality work.
- [heterogeneity:Fail...\(4\)](#)

Minimal Viable Product.

In our environment, Bernie's primary function has been to provide a fallback to Kate's solid work. It has been our playground and our backup server. At a minimum the new server needs to provide an LXD server to test and backup our production containers and virtual machines. As a refence ~~we will~~ *tried* to start at [Rocky Linux's LXD server guide](#).

THE ORIGINALREVISED PLAN.

There are several factors that ~~we won't be able to~~*couldn't have* considered until we ~~are~~*were* actually on the box. Like whether or not the ssd's will work well with the old raid controller. *The installer will find the disks presented by the raid controller. The SSDs were fine, however Rocky is aptly named..*

1. Export images for ernest and the vm's teddy, and franklin.
2. Pull archive disk and mount it's replacement on kb2018.
3. Move teddy(dns2) to kb2018
4. Pull the existing disks from bs2020.
5. Put the ssds into the first two bays and configure the perc to make a single mirrored disk
6. ~~Install rocky linux 8.4 from an iso a dvd or a thumb drive.~~ *Install newest ubuntu release 21.04 (Hirsute Hippo)*
7. BLDGP(3) at https://docs.rockylinux.org/guides/lxd_server/ *the notes from the last lxd server we built combined with a document that I havent written yet*
8. Configure/test disks, lxd, and networking.
9. Copy profiles and images from kb2018
10. Add and configure ansible.
11. Migrate teddy to its new home.

WHAT ACTUALLY HAPPENED.

1. We exported the images and the lxd configuration to the archive disk as described in [ticket: #79](#)
2. We unmounted the archive disk.
3. We shutdown teddy and moved the container to kb2018
4. We pulled and labled all of the disks on bs2020
5. We installed two new 240G SSDs and two new 1TB 10K disks and configured the raid controller to make a raid 1 mirror of the ssds. (the bigger disks will be handled by zfs).
6. We attempted for several hours to install rocky linux 8.4 on the system but could not get the operating system to recognize any installable disks. So, knowing that we would need to adapt whatever changes Ubuntu threw at us in the spring, we installed the newest server release [21.04 \(Hirsute Hippo\)](#)
7. We updated and installed the prerequisites for lxd/w zfs.
8. We configured the network and disks needed to restore the lxd configuration.
9. We restored the lxd configuration from the archive disk.
10. We moved teddy back to the new lxd configuration.
11. We restored the spare containers from the archive disk.
12. Pulled the archive disk off sight.

TODO:

1. Configure ansible on the new server [Ticket #81](#)
2. Update backup scripts to reference new server.[Ticket #82](#)
3. Make backup scripts work with vms [Ticket #58](#)
4. Script archives to create new off site rotating disk.[Ticket #83](#)

REFERENCES

- <https://fatmin.com/2019/11/23/installing-rhel-8-1-on-dell-r710-r610-with-h700-raid-controller/>
- https://docs.rockylinux.org/guides/lxd_server/

FOOTNOTES / SARCASMS

1. Snaps? Juju? Really????
2. See: [trumpery](#)

3. BLDGP/BLGDP = "Build it Like the Dad Gummed Plans". This is a reference to a 70s American Aircraft Modeler editorial on people building tri-planes out of plans for bi-planes and then wondering why they don't fly.
4. *Although thanks to LXD4 we are running Rocky Linux 8.4, along with Centos 7 and Freebsd. So failure is relative.*

2.7 Ansible

2.7.1 susdevadmin/ansible

<https://bitbucket.org/suspectdevicesadmin/ansible/src/master/>

This repository includes the inventory and playbooks for using ansible to manage lxd/lxc based containers as well as common files used to maintain suspect devices environment.

2.7.2 Installing Ansible on lxd server.

[Notes On Installing Ansible](#)

2.7.3 Using Ansible.

/etc/ansible/hosts

The hosts file serves to categorize and document the containers running on kb2018 and bs2020.

Ansible usage/playbook

USING ANSIBLE TO UPDATE CONTAINERS.

```
root@kb2018:/etc/ansible# ansible pets -m raw -a "update.sh"
```

CONTAINER BACKUP/ARCHIVE

Ansible backups are DEPRECATED.

```
cd /etc/ansible ;screen -L time ansible-playbook playbooks/backup-lxd-containers.yml -vvv -i importants
```

A python based script which is run by cron at midnight maintains warm spares on bs2020:

```
/etc/ansible/python/NightlySnapshots.py
```

<https://bitbucket.org/suspectdevicesadmin/ansible/src/master/python/NightlySnapshots.py>

CREATION OF CONTAINERS

Container creation is simply a matter of adding the container info to [/etc/ansible/hosts](#). under local_containers.

```
root@kb2018:/etc/ansible# nano hosts
...
# variables used by playbooks to create/maintain containers
# Host Variables
#   ip_address
#   purpose #
# Lan Variables
#   ip_netmask
#   ip_gateway
#   ip_dns_server (default server)
# Base Image alias to create container
#   image_alias
# Profiles
#   net_and_disk_profile - profile defining disk pool and network connection
#   system_profile      - base users and other cloud-config items
...
[local_containers]
...
agoodauthor ip_address=198.202.31.160 purpose="Sample Server" image_alias="ubuntu/focal/cloud"
...
```

Then run the [create-lxc-containers](#) playbook.

```
root@kb2018:/etc/ansible# ansible-playbook playbooks/create-lxd-containers.yml
```

MAKING YOUR CHANGES PERMINANT

Perminant modifications to this directory should be followed by

```
git commit -a m"Reason for change"  
git push.
```


2.7.4 Notes on installing ansible.

All centralized maintainance should be initiated from a centralized server such as kb2018(colo/spk) or annie (merlot/pdx)

```
root@kb2018:~# apt-get install ansible
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  ieee-data python-asn1crypto python-certifi python-cffi-backend python-chardet python-cryptography python-enum34 python-httplib2 python-idna python-ipaddress
  python-jinja2 python-jmespath python-kerberos python-libcloud python-lockfile python-markupsafe python-netaddr python-openssl python-paramiko python-pkg-resources
  python-pyasn1 python-requests python-selinux python-simplejson python-six python-urllib3 python-xlrd python-xmldict python-yaml
Suggested packages:
  cowsay sshpass python-cryptography-doc python-cryptography-vectors python-enum34-doc python-jinja2-doc python-lockfile-doc ipython python-netaddr-docs
  python-openssl-doc python-openssl-dbg python-gssapi python-setuptools python-socks python-ntlm
...
root@kb2018:~#
```

install python to all containers

```
root@kb2018:~# for h in `lxc list local: -c n --format csv`;do echo $h;lxc exec local:$h -- apt-get install -y python; done
...
root@kb2018:~# for h in `lxc list bs2020: -c n --format csv`;do echo $h;lxc exec bs2020:$h -- apt-get install -y python; done
...
```

seed /etc/ansible/hosts

- localhost (kb2018/annie)

Adding the entry for the localhost is simple

```
root@kb2018:~# nano /etc/ansible/hosts

[pets:children]
servers
containers

[servers]
kb2018    ansible_connection=local
..
root@kb2018:~#
```

- local containers

entries for local containers is equally straightforward.

```
hostname ansible_connection=lxid
```

Which we can generate using lxc list and awk

```
root@kb2018:~# lxc list -c n --format=csv local:|awk '{print $1,"ansible_connection=lxid";}'>>/etc/ansible/hosts
```

- containers on remote host

Containers on the remote host (bs2020) require an additional parameter

```
remotecontainer ansible_connection=lxid ansible_host=remotehost:remotecontainer
```

Which we again generate using lxc list and awk

```
root@kb2018:~# lxc list -c n --format=csv bs2020:|awk '{print $1," ansible_connection=lxid ansible_host=bs2020:"$1;}'>>/etc/ansible/hosts
```

adding access to bs2020 (via ssh to unprivileged account)

Our current security model expressly forbids direct access to all root accounts, users must connect using an ssh key and escalate using their password.

To control a remote server from ansible user (root@kb2018) we:

- create a sudo user for our ansible host

```
root@bs2020:~# useradd kb2018 -c"Governor Kate Brown" -m -g sudo
root@bs2020:~# passwd kb2018
... remember this one for later ...
```

- restrict ssh access to that account to the ip of that particular host.

```
root@bs2020:~# nano /etc/ssh/sshd_config
...
PermitRootLogin no
....
DenyUsers kb2018@"!192.168.31.159,*"
...
root@bs2020:~# service ssh restart
```

- Generate key for our ansible user (root@kb2018)

```
haifisch:~ don$ ssh -p22222 feurig@bs2020.suspectdevices.com
...
Last login: Mon Feb 25 18:56:59 2019 from 97.115.103.251
feurig@kb2018:~$ sudo bash
[sudo] password for feurig:
root@kb2018:~# ssh-key
ssh-keygen ssh-keyscan
root@kb2018:~# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
...
... add ssh key to kb2018@bs2020:~.ssh/authorized_keys ...
...
root@kb2018:~# ssh kb2018@bs2020.suspectdevices.com
```

Testing connectivity.

At this point we can add the remote server to ansible's inventory and check the connectivity.

```
bs2020 ansible_connection=ssh ansible_ssh_user=kb2018
```

"note kb2018 is the localhost, ernest24jan19 (stopped) and douglas are local containers, bs2020 is a remote host and teddy is a container that it hosts"

```
root@kb2018:~# ansible pets -m ping
kb2018 | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
ernest24jan19 | UNREACHABLE! => {
  "changed": false,
  "msg": "... , exited with result 1",
  "unreachable": true
}
...
douglas | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
...
bs2020 | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
...
teddy | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
root@kb2018:~#
```

However we cannot run privileged commands on our remote host.

```
root@kb2018:~# ansible servers -m apt -a "force_apt_get=yes upgrade=yes update_cache=yes autoremove=yes"
bs2020 | FAILED! => {
  "changed": false,
  "msg": "Failed to lock apt for exclusive operation"
}
kb2018 | SUCCESS => {
```

We can fix this by telling ansible to escalate using our user and password

```
bs2020  ansible_host=bs2020.suspectdevices.com ansible_user=kb2018ansible_become=yes ansible_become_user=root ansible_become_pass=my_super_secret_password
```

And we can see that this works. Next we encrypt the password using ansible's vault feature and moving the username and password to the host_vars file.

```
feurig@kb2018:~$ grep 'bs2020 ' /etc/ansible/hosts
bs2020  ansible_host=bs2020.suspectdevices.com ansible_user='{{ bs2020_unpriviledged_user }}' ansible_become=yes ansible_become_user=root
ansible_become_pass='{{ bs2020_become_pass }}'
root@kb2018:~# ansible servers -m apt -a "force_apt_get=yes upgrade=yes update_cache=yes autoremove=yes"
kb2018 | SUCCESS => {
```

"... you are here ..." * create and protect vault password file

```
root@kb2018:~# openssl rand -base64 2048 > /root/.vault_passwd
root@kb2018:~# chmod 600 /root/.vault_passwd
```

- add password file to ansible.cfg

```
root@kb2018:~# nano /etc/ansible/ansible.cfg
..'
# If set, configures the path to the Vault password file as an alternative to
# specifying --vault-password-file on the command line.
#vault_password_file = /path/to/vault_password_file
vault_password_file=/root/.vault_passwd
...
```

- encrypt sudo password

```
root@kb2018:~# ansible-vault encrypt_string 'mybigsecret' --name 'kb2018_become_pass'
kb2018_become_pass: !vault |
    $ANSIBLE_VAULT;1.1;AES256
    663 .... 462
Encryption successful
```

- add user and encrypted password to /etc/ansible/host_vars/bs2020.yml

```
root@kb2018:~# mkdir /etc/ansible/host_vars
root@kb2018:~# nano /etc/ansible/host_vars/bs2020.yml
bs2020_unpriviledged_user: kb2018
bs2020_become_pass: !vault |
    $ANSIBLE_VAULT;1.1;AES256
    663 .... 462
```

- add variables to inventory

```
[pets:children]
servers
containers

[servers]
kb2018  ansible_connection=local
bs2020  ansible_host=bs2020.suspectdevices.com ansible_user='{{ bs2020_unpriviledged_user }}' ansible_become=yes ansible_become_user=root
ansible_become_pass='{{ bs2020_become_pass }}'
#bs2020  ansible_connection=ssh ansible_ssh_user=kb2018

[containers:children]
local_containers
remote_containers

[local_containers]
douglas ansible_connection=lxid
...
[remote_containers]
...
```

```
goethe ansible_connection=lxm ansible_host=bs2020:goethe
```

- and now we can treat all of our pets with the same love and affection.

```
root@kb2018:~# ansible pets -m apt -a "force_apt_get=yes upgrade=yes update_cache=yes autoremove=yes"
```

References/Linkdump

- <https://stackoverflow.com/questions/37297249/how-to-store-ansible-become-pass-in-a-vault-and-how-to-use-it>
- https://docs.ansible.com/ansible/latest/user_guide/vault.html#id6

2.7.5 Python

Using python to directly interact with LXD and Ansible.

Using ansible to create and backup containers is kind of a shitpile (*which is weird because the debian installable pylxd interface looks really good*). Some of the side effects of our attempts to use ansible to back up our containers included snapshots (which had to be manually culled) running at the same ip as the production servers.

For this reason I plan to directly run container creation backup and migration from python (v3). Python scripts should reference the same datasource for which servers are which. My original foray into ansibles internal and poorly documented api felt about as painfull as the ansible itself. I replaced it with `ansible_runner`.

These modules are intended to be usefull standalone or as part of a maintainance routine.

MODULE CONTAINERSHIPDATA

Ansible data

- `ansible_local_containers`
- `ansible_remote_containers`

According to LXD server kb2018

- `kb2018_local_active_containers`
- `kb2018_local_inactive_containers`

According to LXD server bs2020

- `kb2018_remote_active_containers`
- `kb2018_remote_inactive_containers`
- `snapshots`
- `archives`

Cull Lists

- `snapshot_cullist`
- `archives_cullist`

REFERENCES:

- <https://www.programcreek.com/python/example/90872/ansible.parsing.dataloader.DataLoader>
- https://docs.ansible.com/ansible/latest/dev_guide/developing_api.html
- <https://www.programcreek.com/python/example/111311/ansible.inventory.manager.InventoryManager>

2.8 Docker vm configuration

2.8.1 Centos 7 Docker Host (Franklin Rebuild *DRAFT*)

Lxd 4 introduced qemu/vm support making it possible to install docker in a way that doesn't compromise the underlying server. We want to use docker to present a private repository protected by an nginx proxy using LetsEncrypt SSL certificates.

Basic process.

- Install centos7 vm
- Install prerequisites
- Add docker-ce repository
- Install docker-ce
- Install docker-registry
- Install nginx

IN ORDER TO LET DOCKER DO ITS THING WITHOUT LEAKING WE USE A VM.

Create the vm.

```
root@kb2018:/etc/ansible# lxc image copy images:centos/7 local: --copy-aliases --vm
root@kb2018:/home/feurig# lxc init centos/7 franklin --vm -pdefault -psusdev2lvm
```

Add static networking.

```
root@kb2018:/home/feurig# lxc config edit franklin
architecture: x86_64
config:
  image.architecture: amd64
  image.description: Centos 7 amd64 (20210823_07:08)
  image.os: Centos
  image.release: "7"
  image.serial: "20210823_07:08"
  image.type: disk-kvm.img
  image.variant: default
  user.network-config: |
    version: 1
    config:
      - type: physical
        name: eth0
        subnets:
          - type: static
            ipv4: true
            address: 198.202.31.201
            netmask: 255.255.255.128
            gateway: 198.202.31.129
            control: auto
      - type: nameserver
        address: 198.202.31.132
  volatile.base_image: 812cf4c1b46f4ff2422a6e81c9991bdda12b36e53c58f4edc74580e21034860e
  volatile.eth0.host_name: tap62181f92
  volatile.eth0.hwaddr: 00:16:3e:ef:11:ac
  volatile.last_state.power: RUNNING
  volatile.uuid: ce2b1e78-5825-46d1-8335-88caca447a58
  volatile.vsock_id: "50"
devices: {}
ephemeral: false
profiles:
- default
- susdev2lvm
stateful: false
description: ""
root@kb2018:/home/feurig# lxc start franklin
```

AND SINCE THE MACHINE HAD NO NETWORK THE FIRST TIME IT CAME UP IT WON'T HAVE RUN THE CLOUD INIT THAT PROVIDES US WITH OUR USERS USW.

There are centos/7/cloud vms that would let us skip this step.

Install and rerun cloud init.

```

root@kb2018:/home/feurig# lxc exec franklin bash
[root@franklin feurig]# yum install cloud-init
[root@franklin feurig]# cloud-init-cfg all config
[root@franklin ~]# cloud-init clean
[root@franklin ~]# cloud-init init

```

FIRST WE INSTALLED THE DOCKER THAT COMES WITH CENTOS7

```

[root@franklin feurig]# yum search docker
Loaded plugins: fastestmirror, product-id, search-disabled-repos, subscription-
...
Installing:
docker-latest      x86_64      1.13.1-58.git87f2fab.el7.centos      extras      16 M
Installing for dependencies:
criu                x86_64      3.12-2.el7                           base        453 k
docker-client-latest x86_64      1.13.1-58.git87f2fab.el7.centos      extras      3.8 M
libnet              x86_64      1.1.6-7.el7                          base        59 k
protobuf-c          x86_64      1.0.2-3.el7                          base        28 k

Transaction Summary
=====
Install 1 Package (+4 Dependent packages)

Total download size: 21 M
Installed size: 71 M
Is this ok [y/d/N]: y
...
Complete!

```

THEN WE REALIZED ITS TOO OLD AND SO WE GOT THE CURRENT DOCKER-CE FROM DOCKER.**Uninstall what we just did.**

```

[root@franklin feurig]# yum remove docker \
                        docker-client \
                        docker-client-latest \
                        docker-common \
                        docker-latest \
                        docker-latest-logrotate \
                        docker-logrotate \
                        docker-engine
...
Removing:
docker            x86_64      2:1.13.1-208.git7d71120.el7_9      @extras     64 M
docker-client     x86_64      2:1.13.1-208.git7d71120.el7_9      @extras     13 M
docker-client-latest x86_64      1.13.1-58.git87f2fab.el7.centos      @extras     13 M
docker-common     x86_64      2:1.13.1-208.git7d71120.el7_9      @extras     4.4 k
docker-latest     x86_64      1.13.1-58.git87f2fab.el7.centos      @extras     57 M

Transaction Summary
=====
Remove 5 Packages

Installed size: 146 M
Is this ok [y/N]: y
...

```

Add the docker repo and install.

```

[root@franklin feurig]# yum --enablerepo=Extras
[root@franklin feurig]# yum install -y yum-utils
...
Installed:
  yum-utils.noarch 0:1.1.31-54.el7_8

Dependency Installed:
  python-kitchen.noarch 0:1.1.1-5.el7

Complete!
[root@franklin feurig]# yum-config-manager \
  --add-repo \
  https://download.docker.com/linux/centos/docker-ce.repo
...
repo saved to /etc/yum.repos.d/docker-ce.repo

```

Check to see if the docker version is the one that's going to be installed.

```

[root@franklin feurig]# yum list docker-ce --showduplicates | sort -r | head
* updates: mirror.keystealth.org
This system is not registered with an entitlement server. You can use subscription-manager to register.
: manager
Loading mirror speeds from cached hostfile
Loaded plugins: fastestmirror, product-id, search-disabled-repos, subscription-

```

```
* extras: centos-distro.lgservers.com
docker-ce.x86_64          3:20.10.8-3.el7          docker-ce-stable
docker-ce.x86_64          3:20.10.7-3.el7          docker-ce-stable
docker-ce.x86_64          3:20.10.6-3.el7          docker-ce-stable
```

Install it.

```
[root@franklin feurig]# yum install docker-ce
...
Installing:
  docker-ce                x86_64        3:20.10.8-3.el7          docker-ce-stable        23 M
Installing for dependencies:
  containerd.io            x86_64        1.4.9-3.1.el7           docker-ce-stable        30 M
  docker-ce-cli            x86_64        1:20.10.8-3.el7         docker-ce-stable        29 M
  docker-ce-rootless-extras x86_64        20.10.8-3.el7           docker-ce-stable        8.0 M
  docker-scan-plugin       x86_64        0.8.0-3.el7             docker-ce-stable        4.2 M

Transaction Summary
=====
Install 1 Package (+4 Dependent packages)

Total download size: 94 M
Installed size: 380 M
Is this ok [y/d/N]: y
Downloading packages:
warning: /var/cache/yum/x86_64/7/docker-ce-stable/packages/docker-ce-20.10.8-3.el7.x86_64.rpm: Header V4 RSA/SHA512 Signature, key ID 621e9f35: NOKEY
Public key for docker-ce-20.10.8-3.el7.x86_64.rpm is not installed
...
Retrieving key from https://download.docker.com/linux/centos/gpg
Importing GPG key 0x621E9F35:
  Userid      : "Docker Release (CE rpm) <docker@docker.com>"
  Fingerprint: 060a 61c5 1b55 8a7f 742b 77aa c52f eb6b 621e 9f35
  From        : https://download.docker.com/linux/centos/gpg
Is this ok [y/N]: y
...
Installed:
  docker-ce.x86_64 3:20.10.8-3.el7

Dependency Installed:
  containerd.io.x86_64 0:1.4.9-3.1.el7
  docker-ce-cli.x86_64 1:20.10.8-3.el7
  docker-ce-rootless-extras.x86_64 0:20.10.8-3.el7
  docker-scan-plugin.x86_64 0:0.8.0-3.el7

Complete!
```

NOW WE ARE READY TO GO.

```
[root@franklin feurig]# docker run hello-world
docker: Cannot connect to the Docker daemon at unix:///var/run/docker.sock. Is the docker daemon running?.
See 'docker run --help'.
[root@franklin feurig]# systemctl start docker
[root@franklin feurig]# systemctl enable docker
Created symlink from /etc/systemd/system/multi-user.target.wants/docker.service to /usr/lib/systemd/system/docker.service.
[root@franklin feurig]# docker run hello-world
Unable to find image 'hello-world:latest' locally
latest: Pulling from library/hello-world
Digest: sha256:0fe98d7debd9049c50b597ef1f85b7c1e8cc81f59c8d623fcb2250e8bec85b38
Status: Downloaded newer image for hello-world:latest

Hello from Docker!
This message shows that your installation appears to be working correctly.

To generate this message, Docker took the following steps:
 1. The Docker client contacted the Docker daemon.
 2. The Docker daemon pulled the "hello-world" image from the Docker Hub.
    (amd64)
 3. The Docker daemon created a new container from that image which runs the
    executable that produces the output you are currently reading.
 4. The Docker daemon streamed that output to the Docker client, which sent it
    to your terminal.

To try something more ambitious, you can run an Ubuntu container with:
$ docker run -it ubuntu bash

Share images, automate workflows, and more with a free Docker ID:
https://hub.docker.com/

For more examples and ideas, visit:
https://docs.docker.com/get-started/

[root@franklin feurig]#
```

SET UP NGINX AND LET'S ENCRYPT / CERTBOT.**Derp.**

For this we need an fqdn. I picked derp. Derp. Docker Eh? Really? Pfffft.

BLDGP. (Here are some other plans).

The bouncing prompt at <https://www.digitalocean.com/community/tutorials/how-to-secure-nginx-with-let-s-encrypt-on-centos-7> gets us an nginx front end to route our containers through with LetEncrypt ssl certificates that will manage themselves as long as .well-known/acme-challenge is a valid path on the server.

```
[feurig@franklin ~]$ sudo bash
[sudo] password for feurig:
# VVVVV SYSTEM FEEDBACK OMITTED BELOW VVVVV #
[root@franklin feurig]# yum install epel-release
[root@franklin feurig]# yum install certbot-nginx
[root@franklin feurig]# yum install nginx
[root@franklin feurig]# systemctl start nginx
[root@franklin feurig]# systemctl enable nginx
[root@franklin feurig]# nano /etc/nginx/nginx.conf
[root@franklin feurig]# ping derp.suspectdevices.com
[root@franklin feurig]# systemctl reload nginx
```

Cut a hole in the firewall for the http/https server

```
[root@franklin feurig]# iptables -I INPUT -p tcp -m tcp --dport 80 -j ACCEPT
[root@franklin feurig]# iptables -I INPUT -p tcp -m tcp --dport 443 -j ACCEPT
```

And finally install the certificate with certbot.

```
[root@franklin feurig]# certbot --nginx -d derp.suspectdevices.com
Saving debug log to /var/log/letsencrypt/letsencrypt.log
Plugins selected: Authenticator nginx, Installer nginx
Enter email address (used for urgent renewal and security notices)
(Enter 'c' to cancel): don@suspectdevices.com
Starting new HTTPS connection (1): acme-v02.api.letsencrypt.org

- - - - -
Please read the Terms of Service at
https://letsencrypt.org/documents/LE-SA-v1.2-November-15-2017.pdf. You must
agree in order to register with the ACME server. Do you agree?
- - - - -
(Y)es/(N)o: Yes

- - - - -
Would you be willing, once your first certificate is successfully issued, to
share your email address with the Electronic Frontier Foundation, a founding
partner of the Let's Encrypt project and the non-profit organization that
develops Certbot? We'd like to send you email about our work encrypting the web,
EFF news, campaigns, and ways to support digital freedom.
- - - - -
(Y)es/(N)o: Yes
Account registered.
Requesting a certificate for derp.suspectdevices.com and docker.suspectdevices.com
Performing the following challenges:
http-01 challenge for derp.suspectdevices.com
http-01 challenge for docker.suspectdevices.com
Using default addresses 80 and [::]:80 ipv6only=on for authentication.
Waiting for verification...
Cleaning up challenges
Deploying Certificate to VirtualHost /etc/nginx/nginx.conf
Saving debug log to /var/log/letsencrypt/letsencrypt.log
Plugins selected: Authenticator nginx, Installer nginx
Starting new HTTPS connection (1): acme-v02.api.letsencrypt.org
Requesting a certificate for derp.suspectdevices.com
Deploying Certificate to VirtualHost /etc/nginx/nginx.conf
Redirecting all traffic on port 80 to ssl in /etc/nginx/nginx.conf

- - - - -
Congratulations! You have successfully enabled https://derp.suspectdevices.com
- - - - -
Subscribe to the EFF mailing list (email: don@suspectdevices.com).
Starting new HTTPS connection (1): supporters.eff.org

IMPORTANT NOTES:
- Congratulations! Your certificate and chain have been saved at:
  /etc/letsencrypt/live/derp.suspectdevices.com-0001/fullchain.pem
  Your key file has been saved at:
  /etc/letsencrypt/live/derp.suspectdevices.com-0001/privkey.pem
  Your certificate will expire on 2021-11-24. To obtain a new or
  tweaked version of this certificate in the future, simply run
  certbot again with the "certonly" option. To non-interactively
  renew *all* of your certificates, run "certbot renew"
- If you like Certbot, please consider supporting our work by:

  Donating to ISRG / Let's Encrypt: https://letsencrypt.org/donate
  Donating to EFF: https://eff.org/donate-le
```

AND NOW WE ARE READY TO SET UP OUR PRIVATE REPOSITORY IN A DOCKER CONTAINER.

Run up a docker registry:2 image

Change this to use local storage at some point for now the container stores the data

```
[root@franklin feurig]# docker run -d -p 5000:5000 --restart always --name registry registry:2
[root@franklin feurig]# curl localhost:5000/v2/_catalog
{"repositories":[]}
[root@franklin feurig]# docker tag 1fd8e1b0bb7e localhost:5000/registry:2
[root@franklin feurig]# docker push localhost:5000/registry:2
[root@franklin feurig]# curl localhost:5000/v2/_catalog
{"repositories":["registry"]}
```

Configure the Proxy.

What we want is to merge the nginx configuration created by certbot and the one provided below. <https://docs.docker.com/registry/recipes/nginx/> Also the proxy is responsible for authentication.

```
events {
    worker_connections 1024;
}

http {

    upstream docker-registry {
        server registry:5000;
    }

    ## Set a variable to help us decide if we need to add the
    ## 'Docker-Distribution-Api-Version' header.
    ## The registry always sets this header.
    ## In the case of nginx performing auth, the header is unset
    ## since nginx is auth-ing before proxying.
    map $upstream_http_docker_distribution_api_version $docker_distribution_api_version {
        '' 'registry/2.0';
    }

    server {
        listen 443 ssl;
        server_name myregistrydomain.com;

        # SSL
        ssl_certificate /etc/nginx/conf.d/domain.crt;
        ssl_certificate_key /etc/nginx/conf.d/domain.key;

        # Recommendations from https://raymii.org/s/tutorials/Strong_SSL_Security_On_nginx.html
        ssl_protocols TLSv1.1 TLSv1.2;
        ssl_ciphers 'EECDH+AESGCM:EDH+AESGCM:AES256+EECDH:AES256+EDH';
        ssl_prefer_server_ciphers on;
        ssl_session_cache shared:SSL:10m;

        # disable any limits to avoid HTTP 413 for large image uploads
        client_max_body_size 0;

        # required to avoid HTTP 411: see Issue #1486 (https://github.com/moby/moby/issues/1486)
        chunked_transfer_encoding on;

        location /v2/ {
            # Do not allow connections from docker 1.5 and earlier
            # docker pre-1.6.0 did not properly set the user agent on ping, catch "Go *" user agents
            if ($http_user_agent ~ "^(docker\/1\.(3|4|5(?:?!\[0-9]-dev))|Go ).*$" ) {
                return 404;
            }

            # To add basic authentication to v2 use auth_basic setting.
            auth_basic "Registry realm";
            auth_basic_user_file /etc/nginx/conf.d/nginx.htpasswd;

            ## If $docker_distribution_api_version is empty, the header is not added.
            ## See the map directive above where this variable is defined.
            add_header 'Docker-Distribution-Api-Version' $docker_distribution_api_version always;

            proxy_pass http://docker-registry;
            proxy_set_header Host $http_host; # required for docker client's sake
            proxy_set_header X-Real-IP $remote_addr; # pass on real client's IP
            proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
            proxy_set_header X-Forwarded-Proto $scheme;
            proxy_read_timeout 900;
        }
    }
}
```

After installing nginx and running certbot Derp's nginx.conf looks like this.

```
# For more information on configuration, see:
# * Official English Documentation: http://nginx.org/en/docs/
# * Official Russian Documentation: http://nginx.org/ru/docs/

user nginx;
worker_processes auto;
error_log /var/log/nginx/error.log;
pid /run/nginx.pid;

# Load dynamic modules. See /usr/share/doc/nginx/README.dynamic.
include /usr/share/nginx/modules/*.conf;

events {
    worker_connections 1024;
}

http {
    log_format main '$remote_addr - $remote_user [$time_local] "$request" '
        '$status $body_bytes_sent "$http_referer" '
        '"$http_user_agent" "$http_x_forwarded_for"';

    access_log /var/log/nginx/access.log main;

    sendfile        on;
    tcp_nopush      on;
    tcp_nodelay      on;
    keepalive_timeout 65;
    types_hash_max_size 4096;

    include /etc/nginx/mime.types;
    default_type application/octet-stream;

    # Load modular configuration files from the /etc/nginx/conf.d directory.
    # See http://nginx.org/en/docs/nginx_core_module.html#include
    # for more information.
    include /etc/nginx/conf.d/*.conf;

    server {
        if ($host = derp.suspectdevices.com) {
            return 301 https://$host$request_uri;
        } # managed by Certbot

        listen 80;
        listen [::]:80;
        server_name derp.suspectdevices.com;
        return 404; # managed by Certbot

    }
}
```

Adding basic Authentication to the proxy.

But before we add the proxy pass we need to give it some basic authentication. Since we are a 2 admin user system .htpasswd is fine.

```
[root@franklin feurig]# yum install httpd-tools
[root@franklin feurig]# cd /etc/nginx/
[root@franklin nginx]# htpasswd -c .htpasswd feurig
[root@franklin nginx]# nano /etc/nginx/nginx.conf
```

Draft of the nginx.conf

The idea here is to authenticate the docker requests without blocking the certbot parts.

```
# For more information on configuration, see:
# * Official English Documentation: http://nginx.org/en/docs/

user nginx;
worker_processes auto;
error_log /var/log/nginx/error.log;
pid /run/nginx.pid;

# Load dynamic modules. See /usr/share/doc/nginx/README.dynamic.
include /usr/share/nginx/modules/*.conf;

events {
    worker_connections 1024;
}

http {

    ##### >>>>>>> from https://docs.docker.com/registry/recipes/nginx/ >>>>>>>
    upstream docker-registry {
        server localhost:5000;
```

```

}
##### <<<<<< from https://docs.docker.com/registry/recipes/nginx/ <<<<<<

log_format main '$remote_addr - $remote_user [$time_local] "$request" '
                '$status $body_bytes_sent "$http_referer" '
                '"$http_user_agent" "$http_x_forwarded_for"';

access_log /var/log/nginx/access.log main;

sendfile        on;
tcp_nopush      on;
tcp_nodelay     on;
keepalive_timeout 65;
types_hash_max_size 4096;

include /etc/nginx/mime.types;
default_type application/octet-stream;

# Load modular configuration files from the /etc/nginx/conf.d directory.
# See http://nginx.org/en/docs/nginx_core_module.html#include
# for more information.
include /etc/nginx/conf.d/*.conf;

server {
    server_name derp.suspectdevices.com;
    root /usr/share/nginx/html;

    # Load configuration files for the default server block.
    include /etc/nginx/default.d/*.conf;

    error_page 404 /404.html;
    location = /404.html {
    }

    error_page 500 502 503 504 /50x.html;
    location = /50x.html {
    }

    listen [::]:443 ssl ipv6only=on; # managed by Certbot
    listen 443 ssl; # managed by Certbot
    ssl_certificate /etc/letsencrypt/live/derp.suspectdevices.com-0001/fullchain.pem; # managed by Certbot
    ssl_certificate_key /etc/letsencrypt/live/derp.suspectdevices.com-0001/privkey.pem; # managed by Certbot
    include /etc/letsencrypt/options-ssl-nginx.conf; # managed by Certbot
    ssl_dhparam /etc/letsencrypt/ssl-dhparams.pem; # managed by Certbot

    ##### >>>>>> from https://docs.docker.com/registry/recipes/nginx/ >>>>>>
    # disable any limits to avoid HTTP 413 for large image uploads
    client_max_body_size 0;

    # required to avoid HTTP 411: see Issue #1486 (https://github.com/moby/moby/issues/1486)
    chunked_transfer_encoding on;

    location /v2/ {
        # Do not allow connections from docker 1.5 and earlier
        # docker pre-1.6.0 did not properly set the user agent on ping, catch "Go *" user agents
        if ($http_user_agent ~ "(docker\/1\.(3|4|5(?:?!\.([0-9]-dev)))|Go ).*$" ) {
            return 404;
        }

        # To add basic authentication to v2 use auth_basic setting.
        auth_basic "Registry realm";
        auth_basic_user_file /etc/nginx/.htpasswd;

        ## If $docker_distribution_api_version is empty, the header is not added.
        ## See the map directive above where this variable is defined.
        add_header 'Docker-Distribution-API-Version' $docker_distribution_api_version always;

        proxy_pass http://docker-registry;
        proxy_set_header Host $http_host; # required for docker client's sake
        proxy_set_header X-Real-IP $remote_addr; # pass on real client's IP
        proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
        proxy_set_header X-Forwarded-Proto $scheme;
        proxy_read_timeout 900;
    }
    ##### <<<<<< from https://docs.docker.com/registry/recipes/nginx/ <<<<<<
}

server {
    if ($host = derp.suspectdevices.com) {
        return 301 https://$host$request_uri;
    } # managed by Certbot

    listen 80;
    listen [::]:80;
    server_name derp.suspectdevices.com;
    return 404; # managed by Certbot
}

```

```
}}
```

You are here testing/refining this configuration.

REFERENCES.

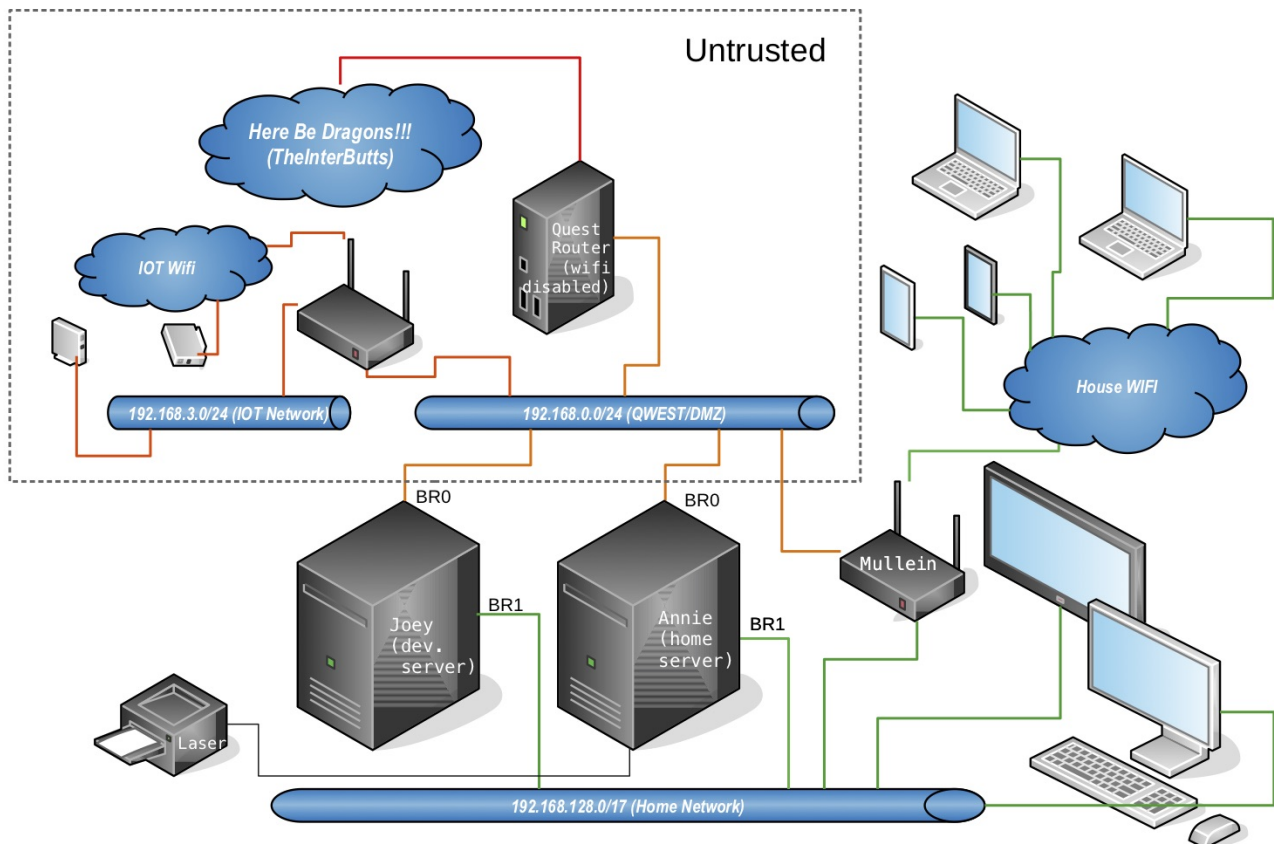
- <https://www.linuxtechi.com/install-docker-on-centos-7/>
- <https://www.linuxtechi.com/setup-docker-private-registry-centos-7-rhel-7/>
- <https://docs.genesys.com/Documentation/System/Current/DDG/InstallationofDockerEngineCommunityEditiononCentOS7>
- <https://blog.simos.info/how-to-use-virtual-machines-in-lxd/>
- <https://www.cyberciti.biz/faq/how-to-secure-nginx-lets-encrypt-on-centos-7/>
- <https://linuxize.com/post/secure-nginx-with-let-s-encrypt-on-centos-7/>
- <https://linuxconcept.com/how-to-secure-nginx-with-lets-encrypt-on-centos-7-linux/>
- <https://stackoverflow.com/questions/41456996/how-to-access-docker-registry-v2-with-curl>
- <https://bobcares.com/blog/docker-private-repository/>
- <https://docs.docker.com/registry/recipes/nginx/>
- <https://www.digitalocean.com/community/tutorials/understanding-nginx-http-proxying-load-balancing-buffering-and-caching>
- <https://serverfault.com/questions/230749/how-to-use-nginx-to-proxy-to-a-host-requiring-authentication>
- <https://www.nginx.com/blog/nginx-plus-authenticate-users/>
- <https://developer.okta.com/blog/2018/08/28/nginx-auth-request>

2.9 Edge server configuration

2.9.1 Joey Rebuild

Joey (ramone) is an edgy server. Joey has 3 main jobs.

1. Stage web content to be pushed or synced to external servers.
2. Provide disk replication for the lan file servers.
3. Provide container space both for lan/wan/edge services and to back up those provided by the file server.



Ubuntu 20.04 + zfs root on the Hp z400.

Someday this will not be so hard :) As much as I like this little workhorse the bios on it kind of sucks. No UEFI. No booting from the on board raid. No booting from the external raid controller.

The current desktop installer can install zfs boot and root disks. It only works with UEFI based bios's but (much like the 18.04 install using the half baked on board psudo-raid controller) *it doesnt notice that your system doesnt support UEFI*. It installs just fine but won't boot. So like the last install I just installed a minimal system and enough zfs tools to detect the installation and let grub find it. (then edit /etc/grub/default to boot to the zfs option and update-grub)

```
... booting from minimal ...
# nano /etc/default/grub
...
GRUB_DEFAULT=2
#GRUB_TIMEOUT_STYLE=hidden
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR=`lsb_release -i -s 2> /dev/null || echo Debian`
GRUB_CMDLINE_LINUX_DEFAULT="verbose"
GRUB_CMDLINE_LINUX="text"
```

```
# update-grub
# reboot
```

While your fixing things get rid of the graphical desktop.

```
... booting from zfs installation ...
root@joey:/# systemctl set-default multi-user.target
```

Then I added the mirror partitions. I didn't bother with the boot partitions on the mirror disk since they didn't work anyway. And I won't bother documenting it because.....

.....FUUUUUUU..... AT THAT POINT I REALIZED I NEEDED TO MAKE SPACE FOR THE CONTAINERS.

Since the installer only accepted a disk name for the zfs install it took the entire disk. Fortunately we are running zfs. Booting from our minimal install we can break the mirrors we just created and then use the extra disks to recreate a smaller boot disk. (Note: root and boot pools were on sde and sdc at this point).

Shrinking a zfs pool.

To shrink the pool size we split the mirror and repartition one of the disks. Then we copy the data to the smaller partition.

```
zpool detach ata-TEAML5Lite3D240G_AB20190109A0101064-part6
zpool export rpool
zpool import rpool oldrpool

mkdir /oldroot /newroot
fdisk /dev/sde
.... delete partition six and split it into 2 new partitions ....
ls -lsa /dev/disk/by-id/|grep sde7
zpool create rpool ata-TEAML5Lite3D240G_AB20190109A0101064-part7
zpool export oldrpool
zpool import -R/oldroot oldrpool
zpool export rpool
zpool import -R/newroot rpool
zfs snapshot -r oldrpool@for_copy
zfs send -R oldrpool@for_copy | zfs recv -F rpool
zpool list
zpool export oldrpool
```

Move boot pool to second disk.

To move the boot pool we attach one the new partitions, mirror it and then detach the original.

```
ls -ls /dev/disk/by-partuuid/|grep sde6
zpool attach bpool ata-TEAML5Lite3D240G_AB20190109A0101064-part6
zpool status bpool
(... wait for resilver to complete ...)
zpool detach bpool d3bff208-06
```

Copy partitions from edited disk.

Once everything is on the newly partitioned disk we can copy the modified partition table to the original disk.

```
sgdisk -p /dev/sde
sgdisk -R/dev/sdc /dev/sde
sgdisk -G /dev/sdc
partprobe
```

Remirror to smaller partitions on original disk.

```
ls -ls /dev/disk/by-partuuid/|grep sdc6
zpool attach bpool ata-TEAML5Lite3D240G_AB20190109A0101064-part6 ata-Crucial_CT240M500SSD1_132909461FE4-part6
ls -ls /dev/disk/by-partuuid/|grep sdc7
zpool attach rpool ata-TEAML5Lite3D240G_AB20190109A0101064-part7 ata-Crucial_CT240M500SSD1_132909461FE4-part7 -o ashift=9
zpool export rpool
zpool import -R/ rpool
zpool status
(... wait for resilvers to complete ...)
update-grub
reboot
```

If I were to do this again.

Now that we know that ubuntu will install a functioning zfs installation I would install the minimal system on the mirror disk and migrate the rpool and bpool to the mirror rather than using a separate disk. I still have that option.

DISK LAYOUT.

- SSDs

id	size	purpose
ata-Crucial_CT240M500SSD1_132909461FE4	240	bpool/rpool/devil
ata-TEAML5Lite3D240G_AB20190109A0101064	240	bpool/rpool/devil
ata-Corsair_Force_GT_1227792800001502028A	120	grub/maintenance disk

- Archive disks

id	size	purpose
ata-Hitachi_HDS5C3030ALA630_MJ1311YNG7RM5A	2.7T	/archive backup
scsi-3600508b1001c407672486f627337a3e9	1.8T	theflatfield/filebox backup
scsi-3600508b1001cca9043287e57e5adae22	1.8T	theflatfield/filebox backup
scsi-3600508b1001cfe22e99aade7378fb6c1	2.7T	/archive backup

NETWORK CONFIGURATION

[etc/netplan/50-cloud-init.yaml](#)

- br0 is the isp router side of the network and provides an anonymous bridge.
- br1 is the internal network and is configured to provide direct connection to the server.

INSTALL SSACLI TO TALK TO THE SAS/RAID CONTROLLER

Like most vendor repositories hp's cant get the signature/otherdata right so we just trust them. (grrrrr)

```
echo deb [trusted=yes] https://downloads.linux.hpe.com/SDR/repo/mcp/ubuntu/ focal current/non-free >>/etc/apt/sources.list
apt-get update
apt-get install ssacli
```

Job #3: Container space (LXD Installation and setup).

```
... snap install lxd --channel=4.0/stable
root@joey:/home/don# apt-get install lxd
...
root@joey:/home/don# lxd init
Would you like to use LXD clustering? (yes/no) [default=no]:
Do you want to configure a new storage pool? (yes/no) [default=yes]:
Name of the new storage pool [default=default]: devil
Name of the storage backend to use (dir, lvm, zfs, ceph, btrfs) [default=zfs]:
Would you like to create a new zfs dataset under rpool/lxd? (yes/no) [default=yes]: no
Create a new ZFS pool? (yes/no) [default=yes]:
Would you like to use an existing empty block device (e.g. a disk or partition)? (yes/no) [default=no]: yes
Path to the existing block device: /dev/disk/by-id/ata-TEAML5Lite3D240G_AB20190109A0101064-part2
Would you like to connect to a MAAS server? (yes/no) [default=no]:
Would you like to create a new local network bridge? (yes/no) [default=yes]: no
Would you like to configure LXD to use an existing bridge or host interface? (yes/no) [default=no]: yes
Name of the existing bridge or host interface: br0
Would you like LXD to be available over the network? (yes/no) [default=no]: yes
Address to bind LXD to (not including port) [default=all]: 192.168.129.65
Port to bind LXD to [default=8443]:
Trust password for new clients:
Again:
Would you like stale cached images to be updated automatically? (yes/no) [default=yes]:
Would you like a YAML "lxd init" preseed to be printed? (yes/no) [default=no]: yes
config:
  core.https_address: 192.168.129.65:8443
  core.trust_password: NOT_HERE
networks: []
storage_pools:
- config:
  source: /dev/disk/by-id/ata-TEAML5Lite3D240G_AB20190109A0101064-part2
  description: ""
  name: devil
  driver: zfs
profiles:
- config: {}
  description: ""
```



```

devices:
  eth0:
    name: eth0
    nictype: bridged
    parent: br0
    type: nic
  root:
    path: /
    pool: devil
    type: disk
  name: default
  cluster: null

```

RESTORING PROFILES AND CONTAINERS FROM LAN FILE SERVER.

```

root@annie:/home/don# lxc remote remove joey
root@annie:/home/don# lxc remote add joey
Certificate fingerprint: 2ad747d1305470bfd6787c1451e0ab64e22ab1798ae78d113718205763639742
ok (y/n)? yes
Admin password for joey:
Client certificate stored at server: joey
root@annie:/home/don# lxc profile copy infra joey:
root@annie:/home/don# lxc profile copy susdev20 joey:
root@annie:/home/don# lxc profile copy susdev21 joey:
root@annie:/home/don# lxc move nina joey:
root@annie:/home/don# lxc start joey:nina

```

ADDING ANSIBLE USER FOR ANNIE.

```

root@joey:# useradd -m -c Annie annie
root@joey:# passwd annie.
...
root@joey:# su -l annie
annie@joey: ~$ ssh-keygen
annie@joey: ~$ nano .ssh/authorized_keys
... paste root@annie public key ...

```

Job #2: Disk replication.

transfer initial large disks from home server.

Smaller disks

Since we are on a private network we can send files in the clear. For small items this only takes a few hours.

- Source Machine

```

root@annie:# zfs snapshot -r filebox@26JAN21
root@annie:# time zfs send -R filebox@26JAN21|pv|nc -l 3333

```

- Destination machine

```

root@joey:# nc annie.local 3333|pv|zfs recv -Fdu filebox

```

Larger disk

The archive disk which has 1.6Tb of data required 30 hours to transfer. I have ordered a pair of jumbo packet capable nics. In theory this should only need to be done once and then deltas can be sent.

```

root@annie:/home/don# time zfs send -R archive@25JAN21|pv|nc -l 3334
1.59TiB 26:04:36 (17.7MiB/s) [
^Iz^Z
(1) Stopped
real    1843m8.561s
user    6m28.958s
sys     96m2.308s
root@annie:/home/don# zfs send -R archive@25JAN21 | pv | nc -l 3334

```

Investigate larger mtu values.

Turns out none of the network adapters on board or cards support jumbo frames.

```
root@annie:/home/don# ip -d link show
...
2: ens6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel master br0 state UP mode DEFAULT group default qlen 1000
    link/ether 00:14:d1:25:2b:bc brd ff:ff:ff:ff:ff:ff promiscuity 1 minmtu 60 maxmtu 7152
...
3: enpls0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master br1 state UP mode DEFAULT group default qlen 1000
    link/ether 78:e7:d1:c3:ef:9e brd ff:ff:ff:ff:ff:ff promiscuity 1 minmtu 60 maxmtu 1500
...
root@joey:/home/don# ip -d link list
...
2: ens1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master br0 state UP mode DEFAULT group default qlen 1000
    link/ether 00:10:18:1b:53:c0 brd ff:ff:ff:ff:ff:ff promiscuity 1 minmtu 60 maxmtu 1500
...
3: enpls0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master br1 state UP mode DEFAULT group default qlen 1000
    link/ether d4:85:64:99:0e:89 brd ff:ff:ff:ff:ff:ff promiscuity 1 minmtu 60 maxmtu 1500
```

Purchased two new jumbo packet capable network cards: one like the [asus nx1101](#). The new cards produced a maximum mtu of 71xx. Given that we add the following to /etc/netplan/50-cloud-init.yaml

```
network:
  version: 2
  renderer: networkd
  ethernets:
    ens6:
      match:
        macaddress: 30:85:a9:38:cc:11
      mtu: 7000
      dhcp4: no
      dhcp6: no
    enpls0:
      dhcp4: no
      dhcp6: no
  bridges:
    br1:
      dhcp4: no
      dhcp6: no
      mtu: 7000
      addresses:
        - 192.168.129.65/17
      gateway4: 192.168.129.1
      nameservers:
        addresses:
          - 192.168.129.1
          - 198.202.31.132
      interfaces:
        - ens6
    br0:
      dhcp4: no
      dhcp6: no
      interfaces:
        - enpls0
```

Job #1: Edgy services.**PIHOLE**

```
apt-get install git-core
git clone --depth 1 https://github.com/pi-hole/pi-hole.git Pi-hole
cd Pi-hole/automated\ install/
bash basic-install.sh
pihole -a -p
pihole enable
service pihole-FTL restart
nano /usr/local/bin/update.sh
... add pihole to update ...
if [ -x "$(command -v pihole)" ]; then
  echo pihole upgrade.
  pihole -up
fi
...
```

SQUID

```
apt-get install squid
nano /etc/squid/squid.conf
...
acl SSL_ports port 443
acl Safe_ports port 80      # http
acl Safe_ports port 21      # ftp
acl Safe_ports port 443     # https
```

```

acl Safe_ports port 70      # gopher
acl Safe_ports port 210     # wais
acl Safe_ports port 1025-65535 # unregistered ports
acl Safe_ports port 280     # http-mgmt
acl Safe_ports port 488     # gss-http
acl Safe_ports port 591     # filemaker
acl Safe_ports port 777     # multiling http
acl CONNECT method CONNECT
http_access deny !Safe_ports
http_access deny CONNECT !SSL_ports
http_access allow localhost manager
http_access deny manager
http_access allow localhost
acl my_internal_net src 192.168.0.0/24
http_access allow my_internal_net
http_port 3128
coredump_dir /var/spool/squid
refresh_pattern ^ftp:      1440  20%  10080
refresh_pattern ^gopher:  1440  0%   1440
refresh_pattern -i (/cgi-bin/|\?) 0    0%   0
refresh_pattern (Release|Packages(.gz)*)$ 0    20%  2880
refresh_pattern .          0    20%  4320

```

WEB SERVER WITH AFS SHARE.

mapping users/drives on the container host

```

echo 'root:1000:1' | sudo tee -a /etc/subuid /etc/subgid
cat /etc/subgid
lxc config set nina raw.idmap 'both 1000 1000'
lxc config edit nina
...
devices:
  sdgl:
    path: theflatfield
    source: /theflatfield
    type: disk
...
lxc start nina

```

Installing netatalk on container

```

apt-get install netatalk
useradd don -m -c "Donald Delmar Davis" -u 1000 -g 1000
passwd don
nano /etc/netatalk/afp.conf
[Global]
; Global server settings
valid users=don
; [Homes]
; basedir regex = /xxxx
[TheFlatField]
path=/theflatfield

```

Serving it up with lighttpd

```

apt-get install lighttpd
nano /etc/lighttpd/lighttpd.conf
...
server.document-root = "/theflatfield/static/digithink/site"
...

```

LINKDUMP

- <https://openzfs.github.io/openzfs-docs/Getting%20Started/Ubuntu/Ubuntu%2020.04%20Root%20on%20ZFS.html#rescuing-using-a-live-cd>
- <https://gist.github.com/yorickdowne/a2a330873b16ebf288d74e87d35bff3e>
- <https://saveriomiroddi.github.io/Installing-Ubuntu-on-a-ZFS-root-with-encryption-and-mirroring/#cloning-the-efi-partition>
- https://www.reddit.com/r/linuxadmin/comments/j8qzdq/install_ubuntu_server_2004_on_a_zfs_root/
- <https://www.medo64.com/2020/04/installing-uefi-zfs-root-on-ubuntu-20-04/>
- <https://serverdocs.suspectdevices.com/serverdocs/wiki/NotesOnAppleTalk3vsUbuntu>
- <https://serverdocs.suspectdevices.com/serverdocs/wiki/TaskInstallSquidCaching>

2.10 Gitea configuration

2.10.1 gitea-configuration

(Build notes for getea server.)

Master Copy: <https://github.com/feurig/gitea-configuration/blob/main/README.md>

Gitea is a github like environment written in go. It provides git in an accessible form and allows you to create issues and write wiki pages like redmine and trac while also serving those repositories.

It is less convoluted than gitlab but more configurable than GCOS which it is based on.

Server Setup

INSTALLING PRE-REQUISITES

We are building on a ubuntu/focal/cloud (from lxc's images) container with preseeded admin accounts.

```
apt-get -y install curl postgresql apache2 git
apt-get install postfix
... add as a Satellite (null client) ...
```

We want to use a single git user so we add it (will deal with this later)

```
adduser --system --shell /bin/bash --group --disabled-password --home /home/git git
```

We are going to use the package provided by packaging.gitlab.io

```
curl -sL -o /etc/apt/trusted.gpg.d/morph027-gitea.asc https://packaging.gitlab.io/gitea/gpg.key
deb [trusted=yes arch=amd64] https://packaging.gitlab.io/gitea gitea main" | sudo tee /etc/apt/sources.list.d/morph027-gitea.list
update.sh
apt-get install gitea
```

SETTING UP POSTGRESQL DATABASE

```
su - postgres
postgres@shelly:~$ createuser -P git
... add passwd
postgres@shelly:~$ createdb gitea -O git
```

INITIAL CONFIGURATION

Once gitea is installed go to myservername:3000 and navigate to the login in the upper right corner. Fill in the database,username, and dbpassword. Replace localhost with your servers fqdn. Create admin user (remember password here)

Testing it out.

The first thing we want to do here is to mirror our github repositories.

MIRRORING GITHUB REPOSITORIES.

We want to automate mirroring all of our repositories hosted on github (and bitbucket at some point). To do this we create a personal-access-token from our github developer tools. (save the token somewhere as it will not be recoverable). Once we have that token we select New migration. Fill in the <https://github.com/myuser/myrepo> and paste the token into the form, select mirror and the magic begins.

EDITING THE MIRROR INTERVAL

The default mirror interval is 8 hours with a minimum of 10 minutes. To fix this we add the following to `/etc/gitea/app.ini`

```
nano /etc/gitea.app.ini
...
[cron.update_mirrors]
SCHEDULE = @every 2m
```

```
[mirror]
DEFAULT_INTERVAL = 1h
MIN_INTERVAL = 2m
...
service gitea restart
```

AUTOMATING CREATION OF MIRRORS (GITHUB).

We were able to automate mirroring our github repos with the help of some python provided by jpmens.net we modified it to allow us to separate local mirrors by the same organizations used by github though this required us to manually add the local users and organizations. The script in progress is here.

<https://github.com/feurig/gitea-configuration/blob/main/mirror-repos.py>

Manually migrating bitbucket mirrors.

Like github mirroring bitbucket repositories required the creation of an application password. Then add a new "Git" migration using the application password as your credentials.

Setting up ssl and apache proxy.

Gitea runs as an unprivileged user on port 3000. To present it as a normal web server required a proxy server (apache). Since we had created letsEncrypt certificates for the old git server we moved them. There are still permissions issues with giving gitea access to the certificates which were worked around by copying the files.

Getting apache to proxy the https required enabling 'proxy_http2' and 'proxy' module (not 'proxy_http')

```
a2enmod proxy proxy_http2
```

- Gitea configuration [etc/gitea/app.ini](#)
- Apache configuration [etc/apache2/sites-available/gitea.conf](#)

MANUALLY UPDATING LETSENCRYPT CERTIFICATES.

Gitea serves static content under the public/custom directory. In order to update the lets encrypt certificates you will need to open two shells into the git server.

In the first window initiate the update request.

```
certbot certonly --manual
Saving debug log to /var/log/letsencrypt/letsencrypt.log
Plugins selected: Authenticator manual, Installer None
Account registered.
Please enter in your domain name(s) (comma and/or space separated) (Enter 'c'
to cancel): git.suspectdevices.com
Generating a certificate request for git.suspectdevices.com
Performing the following challenges:
http-01 challenge for git.suspectdevices.com

- - - - -
Create a file containing just this data:

loIND53_1yZLSZX1lKkLqhCBY7YhNo_vzdyrEznXHSQ.MQ0FIIV4g1TA8EJatJpiciDipeqHIMHJsetBrs2tzqM

And make it available on your web server at this URL:

http://git.suspectdevices.com/.well-known/acme-challenge/loIND53_1yZLSZX1lKkLqhCBY7YhNo_vzdyrEznXHSQ

- - - - -
Press Enter to Continue
```

In the second window create the file requested file.

```
root@shelly:/var/lib/gitea# cd public/custom/
root@shelly:/var/lib/gitea/public/custom# echo loIND53_1yZLSZX1lKkLqhCBY7YhNo_vzdyrEznXHSQ.MQ0FIIV4g1TA8EJatJpiciDipeqHIMHJsetBrs2tzqM >.well-known/acme-challenge/
loIND53_1yZLSZX1lKkLqhCBY7YhNo_vzdyrEznXHSQ
root@shelly:/var/lib/gitea/public/custom# chown -R gitea:gitea .
```

Then continue in the first window and copy the new keys to where gitea expects them.

```
Press Enter to Continue
... when finished copy the new certs to gitea ...
cd /etc/letsencrypt/live/git.suspectdevices.com/
```

```
cp fullchain.pem privkey.pem /var/lib/gitea/keys/  
chown -R gitea:gitea /var/lib/gitea/keys/  
reboot
```

TODO (NO MAJOR ISSUES)

- Document making gitea less ugly (add Susdev brand look and feel)
- Fix permission issues with certificate issues to allow for autorenewal if possible.
- Consider normalizing git user and repo locations.

references/linkdump

- <https://gitlab.com/packaging/gitea>
- <https://bryangilbert.com/post/devops/how-to-setup-gitea-ubuntu/>
- <https://luxagraf.net/src/gitea-nginx-postgresql-ubuntu-1804>
- <https://docs.github.com/en/free-pro-team@latest/github/authenticating-to-github/creating-a-personal-access-token>
- <https://jpmens.net/2019/04/15/i-mirror-my-github-repositories-to-gitea/>
- <https://websiteforstudents.com/how-to-install-gitea-git-server-on-ubuntu-16-04-18-04-18-10-with-mariadb/>
- <https://docs.gitea.io/en-us/config-cheat-sheet/>
- <https://charlesreid1.github.io/setting-up-a-self-hosted-github-clone-with-gitea.html>
- https://charlesreid1.com/wiki/Gitea#Using_Binary
- <https://mindefrag.net/2018/07/how-to-install-and-configure-gitea-a-self-hosted-github-like-service/>

2.11 Mkdocs server configuration

2.11.1 Mkdocs Server Configuration.

For the past several years I have been using Trac to maintain server notes and create a todo sort of ticketing system for our systems. Trac is kind of a pain to setup and maintain and though I work well with the wiki (documenting as I go), and the ticketing communicates the work being done, I have not been able to get others to contribute to the documentation. I tried the trac extension to allow markdown to be embedded in the wiki pages but its kind of janky.

Recently I started documenting my builds and projects in markdown and then storing them along with any configuration files and scripts in Github or Bitbucket repositories. Assuming that this is the way forward I am rebuilding the digithink.com site using markdown as the source.

Markdown based web services.

After looking at several options I narrowed my search to mk-docs and allmark. Mkdocs was in the supported ubuntu repos so I started there. I got good results but was disappointed with the themeing available, until I looked at [Material for Mk-docs \(https://squidfunk.github.io/mkdocs-material/\)](https://squidfunk.github.io/mkdocs-material/)

INSTALLING MATERIAL FOR MK-DOC

mkdocs-material unfortunately isn't packaged however, it can be installed through pip3. To export the complete site as a pdf we also install the mkdocs-with-pdf.

```
apt-get install python3-pip
pip3 install mkdocs-material
apt-get install build-essential python3-dev python3-pip python3-setuptools python3-wheel python3-cffi libcairo2 libpango-1.0 libpangocairo-1.0 libgdk-pixbuf2.0
libffi-dev shared-mime-info
pip3 install WeasyPrint
pip3 install mkdocs-with-pdf
```

CONVERTING TRAC WIKI ENTRIES TO MARKDOWN

I was able to convert the 50 or so wiki pages on serverdocs and clean them up.

Ruby

I found a [gist \(and three refinements\)](#), which even though I don't ruby well I was able to adapt to pg-ruby. It isn't perfect but it worked.

[export.rb](#)

Python

I found this [python based script \(trac2down.py\)](#) which needed to be adapted for postgres and python 3. It does not handle tables. It does however insert an author and timestamp into the document. If I can I would like to finish this with table support. In the mean time, I have a working copy that is clean enough.

[export.py](#)

... add some explanation for getting the resulting the converted markdown to the mkdocs servers...

BUILDING THE STATIC WEB SITE (STAGING SERVER).

```
cd /theflatfield/static/digithink/
mkdocs build
chown -R www-data:www-data site/
```

SERVING IT UP WITH LIGHTTPD

```
root@nina:/# nano /etc/lighttpd/lighttpd.conf
...
server.document-root    = "/theflatfield/static/digithink/site"
server.upload-dirs       = ( "/var/cache/lighttpd/uploads" )
server.errorlog          = "/var/log/lighttpd/error.log"
server.pid-file          = "/run/lighttpd.pid"
```

```

server.username      = "www-data"
server.groupname     = "www-data"
server.port          = 80
...
root@nina:/# service lighttpd restart

```

MOVING THE SITE TO PRODUCTION.

The source code for the site is on github under feurig/digithink. Any local changes made to the source files (like this one) should be pushed.

```

hoffa:docs don$ git commit -a -m "Start documenting process for mkdoc -> static site"
...
hoffa:docs don$ git push

```

Then the site is updated and rebuilt on herbert our lighttpd server (serves digitihiink,busholini, and 3dangst).

```

root@kurt:~# cd /var/www/digithink
root@kurt:/var/www/digithink# git pull
remote: Enumerating objects: 21, done.
remote: Counting objects: 100% (21/21), done.
remote: Compressing objects: 100% (10/10), done.
remote: Total 16 (delta 5), reused 15 (delta 4), pack-reused 0
Unpacking objects: 100% (16/16), done.
From github.com:feurig/digithink
   89bf97d..d563f50  main       -> origin/main
Updating 89bf97d..d563f50
Fast-forward
 .gitignore                  | 3 ++
 docs/buildnotes/mkdocs-server-configuration/export.py              | 68 +++++
 .../mkdocs-server-configuration}/export.rb                        | 0
 .../mkdocs-server-configuration.md                                | 56 +++++
 4 files changed, 127 insertions(+)
 create mode 100644 .gitignore
 create mode 100644 docs/buildnotes/mkdocs-server-configuration/export.py
 rename docs/{legacy/serverdocs => buildnotes/mkdocs-server-configuration}/export.rb (100%)
 create mode 100644 docs/buildnotes/mkdocs-server-configuration/mkdocs-server-configuration.md

```

Since we created several buildnotes repositories before we started this project we imported them as submodules. We need to sync and update these as well.

```

root@kurt:/var/www/digithink# git submodule sync
Synchronizing submodule url for 'docs/buildnotes/edge-server-configuration'
Synchronizing submodule url for 'docs/buildnotes/gitea-configuration'
Synchronizing submodule url for 'docs/buildnotes/redmine-configuration'
root@kurt:/var/www/digithink# git submodule update

```

Then we update the live site.

```

root@kurt:/var/www/digithink# mkdocs build
INFO - Cleaning site directory
INFO - Building documentation to directory: /var/www/digithink/site
INFO - Documentation built in 3.22 seconds
root@kurt:/var/www/digithink# chown -R www-data:www-data site/

```


2.12 Redmine configuration

2.12.1 Redmine Server.

<https://github.com/feurig/redmine-configuration/blob/main/README.md>



Suspect Devices maintains a git backup server for repositories hosted by github and bitbucket. This site uses Redmine to track issues and work.

TASKS

- Backup repositories hosted elsewhere.
- Consolidate work into active/inactive projects
- Track issues (ticketing)
- Document server setup.

2.12.2 Server configuration

This server is running on a Ubuntu 18.04 container because redmine requires a version of Ruby that is behind the new LTS (20.04). We will revisit this next spring.

```
apt-get install postgresql
apt-get install apache2 libapache2-mod-passenger
apt-get install redmine-pgsql
apt-get install redmine
cp /usr/share/doc/redmine/examples/apache2-passenger-host.conf /etc/apache2/sites-available/redmine.conf
nano /etc/apache2/sites-available/redmine.conf
a2enmod passenger
a2ensite redmine.conf
a2dissite 000-default
service apache2 reload
update.sh
```

Adding git functionality...

```
apt-get install git
```

Add git command to configuration

```
cp /usr/share/redmine/config/configuration.yml.example /etc/redmine/default/configuration.yml
nano /etc/redmine/default/configuration.yml
... add git command here ...
  scm_git_command: git
...
service redmine restart
```

Create some space for mirrors.

```
mkdir /var/git
chown -R www-data:www-data /var/git/
```

The www-data user should have its keys added to bitbucket and github. (This user does not need write permission)

```
vipw
su - www-data
mkdir /var/www/.ssh
chown www-data:www-data /var/www
su - www-data
ssh-genkey
```

Rather than configuring a git hook for both github and bitbucket we will create scripts to populate and update the mirrors.

```
vi /etc/cron.d/sync_git_repos
*/2 * * * * www-data /var/www/bin/update-repos.py
```

MAKING REDMINE LESS UGLY.

Redmine makes it fairly easy to theme using css to override its defaults.

```
cd /usr/share/redmine/public/themes/
ls
mkdir susdev
chown www-data susdev
mkdir susdev/stylesheets/
mkdir susdev/images
ls
cd susdev/images/
wget https://serverdocs.suspectdevices.com/serverdocs/chrome/site/sd_logo_sm.png
wget https://serverdocs.suspectdevices.com/serverdocs/chrome/site/sd_logo_sm.png --no-check-certificate
nano ../stylesheets/application.css
ls
nano ../stylesheets/application.css
chown -R www-data:www-data ../../susdev
```

[usr/share/redmine/public/themes/susdev/stylesheets/application.css](#)

ADDING SSL TO THE SITE

```
sudo bash
make-ssl-cert generate-default-snakeoil --force-overwrite
cd /etc/apache2/
ls
a2enmod ssl
nano sites-enabled/redmine.conf
apache2ctl configtest
apache2ctl restart
```

Getting a certificate from letsencrypt

the EFF provides a certificate and a program to set it up from letsencrypt

```
apt-get install certbot
```

Certbot expects to be able to verify that your server exists and can serve one of its files. The file needs to be accessible at `http://\well-known/acme-challenge/` the example below assumes the document root for redmine.

```
cd /usr/share/redmine/public
mkdir -p .well-known/acme-challenge/
echo hello> .well-known/acme-challenge/test
root@emile:/usr/share/redmine/public# chown -R www-data:www-data .well-known/
```

Once this is done you can run certbot manually.

```
certbot certonly --manual
```

They are going to ask a bunch of questions and then ask you to create file on the server. The script pauses and you will have to create the file in a different shell.

```
- - - - -
Create a file containing just this data:

KncX49YdVo125HQZiI1qYbSZxIPIUPMmcJUg2thHHCs.yo0bxAOItnb_LvbpT7eC0ZwNmD_R0uCOAkQqFAoKSTc

And make it available on your web server at this URL:

http://git.suspectdevices.com/.well-known/acme-challenge/KncX49YdVo125HQZiI1qYbSZxIPIUPMmcJUg2thHHCs

- - - - -
Press Enter to Continue
```

Create the file as instructed in a different terminal and make sure its accessible by apache.

```
echo KncX49YdVo125HQZiI1qYbSZxIPIUPMmcJUg2thHHCs.yo0bxAOItnb_LvbpT7eC0ZwNmD_R0uCOAkQqFAoKSTc>/usr/share/redmine/public/.well-known/acme-challenge/
KncX49YdVo125HQZiI1qYbSZxIPIUPMmcJUg2thHHCs
chown www-data:www-data /usr/share/redmine/public/.well-known/acme-challenge/KncX49YdVo125HQZiI1qYbSZxIPIUPMmcJUg2thHHCs
```

If it's successful it will install the certificate and private key under `/etc/letsencrypt/live/`. Adjust your apache configuration.

```
nano /etc/apache2/sites-enabled/redmine.conf
... replace the top portion of the original virtualhost config with the following ....
<VirtualHost *:80>
Redirect permanent "/" "https://git.suspectdevices.com/"
</VirtualHost>

<VirtualHost *:443>

    ServerName git.suspectdevices.com
    SSLEngine on
    #SSLCertificateFile      /etc/ssl/certs/ssl-cert-snakeoil.pem
    #SSLCertificateKeyFile   /etc/ssl/private/ssl-cert-snakeoil.key
    SSLCertificateFile       /etc/letsencrypt/live/git.suspectdevices.com/fullchain.pem
    SSLCertificateKeyFile    /etc/letsencrypt/live/git.suspectdevices.com/privkey.pem

    # this is the passenger config

... and save it ....
apache2ctl configtest
apache2ctl restart
```

- [etc/apache2/sites-enabled/redmine.conf](#)

CREATING SCRIPTS CLONE AND UPDATE THE REPOSITORIES

both bitbucket and git have apis that allow you to list the repositories for each user without needing to authenticate (and expose your credentials). There are limitations but they are worth exploring.

```
apt-get install python-github
apt-get install python-bitbucket

su -l www-data
python
```

The scripts I arrived at work but could certainly be refined. I should probably just use a list for each repo regardless of the site and maintain that as part of this repo. Bitbucket does not allow you to list all of the private repos so I just went with a simple list.

- [clone-repos.py](#)
- [update-repos.py](#)

SET UP EMAIL

Debian's postfix installer makes it very easy to install postfix configured as a null client. When installing select Satellite and provide your domain name and relay host.

```
apt-get install postfix
```

Things that are done in redmine.

- Set passwords and add admin users.
- Add projects and add repositories to them.
- Remove repo browsing from anonymous / non project users.
- Activate theme.
- USE IT!

3. Resume

3.1 D Delmar Davis

Portland, Oregon, ddelmardavis@gmail.com (503) 284-2945

3.1.1 Summary

(I'll be 78 when Unix(tm) time ends)...

I have been administering Unix systems for more than 30 years, in addition to deploying and maintaining Web and other Internet services for 25. The systems have ranged from stand-alone, completely exposed servers to services with separate Web, application, and database layers, clustered and load balanced for redundancy and scalability, and placed behind firewalls, I have done work in the ever amorphous cloud infrastructure provided by Amazon but I am more interested in locally owned and operated LXD based containers.

I have extensive experience with Linux deployments (Red Hat/Fedora/Centos, SUSE, Debian/Ubuntu, Tizen, OpenWRT) I have professional experience with HP-UX, True64, Linux, Solaris, FreeBSD, OS X, and AIX and I have installed and configured databases such as Oracle, Sybase, and SQL Server. Against my better judgment, I have installed, maintained and configured Microsoft systems, and even made them play well with others (SSO, SMB, etc). Additionally, I have a strong background and proven success in providing instruction and documentation for work to be maintained by others.

Throughout my career I have learned new systems quickly. I follow issues and problems through to solution, be they social or technical. Security is the key to surviving on both the World Wide Web and large Intranets. For this reason, I am a practitioner and proponent and of well defined policies, best practices, regular updates, and common sense. My reputation for getting things done and team participation make me an ideal candidate for positions where honest work is valued. Specialties: Thinking outside the box.

Also: Springfield and Waltham are not Boston in the same way Hillsboro and Beaverton are not Portland and I am not interested in your urgent requirement for Intel's permanently contingent workforce at Jones Farm so don't bother asking....

3.1.2 Experience

Systems Engineer

Laika Jul 2021 - Present.

Contiguous Online Server Presence

Fromhell.com Nov 1996 - Present (*24 years 7 months +*)

I have built and maintained servers for the domains digithink.com and fromhell.com (email only) since 1996. They started as a sparc IPC running SunOs 4.1.3 and have had many different variants of BSD and Linux since then. They are currently on LXD containers running Ubuntu (20.04 LTS). My current toolset includes Ansible, Python and ZFS. I also document my work even when it's just mine. (See: <https://www.digithink.com>)

Package Handler

UPS Oct 2017 - Jul 2021 (*3 years 9 months*)

It's funny, If I went to the gym I would never build the muscles I have from sorting ~6000 packages a day for close to 4 years. I had a 12 minute bicycle commute. Another 10 minutes to get to the other side of the channel to PCC's campus in the shipyards (Before COVID, UPS reimbursed my tuition as I updated my welding skills). I got to see first hand the advantages that unionized workers have over similar non union positions (looking at you Amazon). Through the Teamsters I received better health insurance than I ever received doing tech work. Also nice to not have a digital leash (going on 3 years cell phone free). Life is good.

Applications Engineer

Suspect Devices Apr 2008 - Jan 2017 (*8 years 10 months*)

For roughly a decade I worked through Tempus Dictum (DBA Suspect Devices) building hardware and software solutions around the Arduino and other open source hardware and software. My primary focus was on tools for artists and musicians. Client projects ranged from medical equipment, to bicycle racing, to music, to heavy industrial motor controllers. In addition to my embedded work I coded applications for Macintosh and iPhone. I developed hardware platform for teaching micro-controllers to artists and hobbyists. I produced and continue to present workshops focused on introducing micro-controllers to the community. As a systems administrator I maintained multiple internet servers for this company and clients. Upgraded and maintaining a linux based scientific computing cluster at UCSF. Test-deployed cluster implementation to the Amazon Elastic Cloud to benchmark cloud against old hardware. Recent client work involved moving a legacy FreeBSD System to the cloud and exploring cloud based archival options.

IT Administration Engineer

Jaguar Land Rover Sep 2013 - Aug 2015 (*2 years*)

I was contracted to provide primary IT support for the servers and workstations at the JLRNA's Open Software Technology Center in Portland. Supported the build servers and developers tools for the tizen platform as prescribed by Intel (OpenSuse 12.1-12.3). Later I was hired. I migrated servers to new Open Source Technology Center in the Pearl and argued successfully to move to Debian as primary linux platform. I was unsuccessful in convincing management that the position should be outsourced.. Supported 56 servers running Debian and OpenSUSE in addition to roughly the same number of Windows 7 based systems and users.

Continuing Education Instructor

Pacific Northwest College of Art Sep 2012 - Dec 2014 (*2 years 4 months*)

After three years of giving one day classes to the community through Dorkbotpdx I was asked by PNCA's extensions school to develop curriculum for and teach an 8 week survey course introducing Artists to microcontrollers using the Arduino platform.

Unix Systems Engineer

Adecco Dec 2012 - Mar 2013 (*4 months*)

I fulfilled a 3 month contract with Integra focused on covering staff shortfalls (2 people covering ~170 business critical systems). This work included a security audit of key systems, ongoing Oracle upgrade support, as well as replacing much of the expensive and complicated CA Spectrum suit with Nagios and other open source monitoring tools. Worked with senior administrator to streamline and clean up administration. Helped to evaluate free version of puppet versus CF engine and hand rolled scripting. And in general kept the rubber side down.

Senior Network Analyst

Washington County Mar 2007 - Nov 2007 (*9 months*)

Built and configured systems for Oracle RAC cluster using SLES on generic SAN connected blades to replace expensive L Class HP Servers. Moved all production data from HP-UX file server to Novell OES server on SAN connected blade. (Reducing Costs / Increasing Performance).

Unix System Administrator

Washington County May 2006 - Dec 2006 (*8 months*)

Tested and documented file restoration on business critical systems as well as creating process and media for bare bones AIX recovery. Developed transition plan from aging HP-UX servers to san attached generic blade architecture. Deployed first workstation in this process. Worked with IT Services (ITS) and one of its vendors to resolve several support issues. Interviewed, helped select, and trained permanent Unix System Administrator. Created operations and troubleshooting guides for all ITS Unix systems. Worked toward better integration between Unix SA position and ITS support team. Upgraded mediawiki server and

trained ITS staff to use it for its internal documentation. Worked with application and ITS teams to provide additional san space, file restoration, additional printers and modems as needed.

Unix Consultant

SolutionsIQ May 2005 - May 2006 (*1 year 1 month*)

Prepared Unix production and development systems for maintenance by windows based skeleton crew. Systems were comprised of Linux and Solaris based Oracle servers along with several Debian- based infrastructure servers. Cloned Solaris 8/Oracle 8 server for disaster recovery. Spec'd out and installed additional disk for all Solaris Oracle servers. Built serial console server for sun systems. Trained operations staff in general maintenance tasks for all Unix and Linux platforms.

Unix System Administrator

SolutionsIQ Sep 2005 - Dec 2005 (*4 months*)

Researched, redesigned and deployed web server and content management framework for internal content and documentation. Integrated existing documentation into framework. Trained team to perform ongoing maintenance. Deployed 11 production and development servers running Solaris 9 and 10. Worked with other engineers to refine Jumpstart install process, in particular with regards to JASS and postinstall scripts.

Network Engineer

Hewlett Packard Enterprise Sep 2002 - May 2005 (*2 years 9 months*)

Maintained all aspects of 150 system proprietary development environment for HP's internal web development at remote data center. Managed migration of initial environment to more stable data center. Coordinated the addition 70 servers to meet expanded capacity needs. Worked with other administrators to keep all systems patched in response to constant security updates (over 400 systems, HP-UX/Linux/W2K). Built 4 Terabyte HP-UX/Samba solution for Windows cluster suffering unplanned exponential storage growth. Built and maintained HP-UX build server for Linux distributions. Helped develop organizational security policy. Provided 24x7 tier-1 support for live applications as member of 2 teams spanning 5 data centers. Worked with development teams to prepare troubleshooting guides clear enough to outsource support. Initiated password audit to bring production systems in line with security policy. Maintained professional level of service and support to organization which suffered 7 re- organizations over a two and a half year period.

UNIX System Administrator

Rogue Wave Sep 2001 - Apr 2002 (*8 months*)

Worked as a member of 5-person team maintaining over 90 UNIX servers used to develop, build, and troubleshoot Rogue Wave's C++ library products. Installed and configured 30 systems running Solaris, AIX, HP-UX, True64 and Linux. Primarily responsible for new Solaris and AIX build servers. Installed and configured Oracle and Sybase databases. Worked with Senior Admin to develop centralized, scripted system setup to allow rapid deployment/ recovery of system configurations for a given software release. Implemented / maintained configuration scripts for Solaris (2.6 - 9beta) and AIX (4.33 - 5.1) systems. Migrated production NIS and license servers from arcane and dying systems to supportable hardware and OS levels. Set up demo server for new web-based technology. Installed compilers, databases, patches, and other software required for development.

Fabricator

GIBSON STEEL FABRICATING, INC. Jun 2001 - Sep 2001 (*4 months*)

I told my employer that if they called me at 11am on my days off, as they had for months, I was walking. You have to mean that stuff. To be fair it took well over a month before I got the call, gave my notice, and went back to welding. I stood and welded out catch basins (GMAW and OXY/Acetylene work) until I realized how short I was financially. I asked for \$1.75 more an hour at a shop that hadn't raised any of its employees wages in close to a year. They raised the entire shop floor by 50c and offered me a 75c raise. I took the first tech job offered. Like I said. You have to mean it.

UNIX System Administrator / Database Programmer**Modern Medium** Aug 2000 - Jul 2001 (*1 year*)

Systems/Database Administrator and Programmer for www.buymusichere.com, a stocked 250,000 product virtual storefronts for 15 clients with 100 projected. Previous contractor was unable to produce more than 4 working storefronts in 1 year; in 3 months we produced 15 storefronts. Eliminated redundant data, reducing processing time and storage needs by ~70%. Established backup of databases and system. Reconfigured raid system to use existing resources effectively. Installed additional memory, disk and processor to allow for growth. Set up development environment on smaller Sun. Wrote import scripts in Transact SQL/PHP for the automatic updating of databases.

President**Digithink** Jun 1996 - Dec 2000 (*4 years 7 months*)

Set up a small ISP while attending college. Consulted on various jobs, resolved Sun hardware/DNS issues for Ordata.com (now Willamette.net). Net presence provided the basis for contract work and fiscal stability between major contracts/

UNIX System Administrator / Programmer**Northwest Media** May 2000 - Aug 2000 (*4 months*)

Created development environment for web site geared towards post care tracking of youth from programs such as foster care, jobs plus, and job core. Set up UNIX (FreeBSD) based development environment using Staging / File server behind a firewall. Ported web-server based data entry to user-friendly firewall protected Access/VBA application, wrote DLL's to publish data to server.

UNIX Systems Specialist/Database Programmer**Oregon Public Education Network** Feb 1997 - Jun 1999 (*2 years 5 months*)

Provided system administration, programming, and technical consulting to the OPEN-C website [http:// www.open.k12.or.us](http://www.open.k12.or.us). Deployed three web servers with systems running Solaris and HP-UX (everything from boxes of parts to web sites). Developed SQL/WWW scripting language in Perl, consolidating most of sites cgi-scripts. Later contracted to redesign system as an Apache Module. Provided technical consulting, diagnosing and resolving all UNIX Network and World Wide Web related issues.

System Administrator and Security Consultant

DNSI Sep 1996 - Jul 1997 (*11 months*) Instituted security audit of systems and made recommendations to improve operational security. Proposed and implemented plan to restructure LLC company with financial debt nearing \$18,000 which incorporated upgrading Internet connectivity while reducing costs, and restructuring company to provide financial solvency and stability. Served as liaison with US West, guiding company through complex series of phone line shortages. Arranged transfer of hardware and software for 700 clients to new servers and location. Daily operations included setup and administration of FreeBSD and Linux servers from building/ installing hardware, OS, services, and virtual hosts all the way through client relations.

SHOP HAND/FABRICATOR**GIBSON STEEL FABRICATING, INC.** Jul 1995 - Sep 1996 (*1 year 3 months*)

When I moved to Eugene the wage base was so low that Symatec moved its customer support there. So I went to work using the skills I learned in High School (Welding). I performed all aspects of storm water catch basin assembly except for seam welding. Operated hydraulic sheers, torch ,band saw, and fork lift. Assembled basins using SMAW.

Technician**Eli Hefron and Sons** Jan 1993 - May 1993 (*5 months*)

Tested setup and configured surplus sun systems, including installation of SunOS 4.1.x through Solaris 1.x [sic] for shipment to clients.

UNIX System Administrator/CAD Support Specialist**Badger Engineers** Feb 1989 - Sep 1992 (*2 years 7 months*)

Supported department's expansion from one Vax and 4 un-networked PCs to two Vax's, 40 Unix Stations, and 100 networked PCs. UNIX administration included direct user support, adding software, user accounts and scripting. Served as project leader; responsible for development of a method of high volume batch translation between different CAD formats. Piloted the use of several (then) new technologies for better cross platform integration such as PCNFS and sendmail based problem logging. Initiated cadre based help/support groups. Trained senior operators and engineers in computer fundamentals.

Engineering Programmer

Bovay Northwest, Inc. 1986 - 1988 (3 years) Automated drafting and design processes; programming in AutoLisp, MuLisp, and DBase.

3.1.3 Education

University of Oregon**Bachelor of Arts, History** 1995 - 2004

At UO I studied History with an emphasis on revolution. While there I actively promoted punk and other local music as a DJ at KWVA radio. I was invited to display my artwork in several solo and group exhibitions.

Portland Community College**Career Pathways Certificate, Welding Technology/Welder** 2017 - 2017**Spokane Falls Community College**

Associate of Arts, Science; Software Engineering Technology 1983 - 1991 At SFCC I split my studies between fine arts, core studies and software engineering.

Kellogg Sr High**High School Diploma, High School** 1982 - 1983

Cross country, debate, yawn.

Secondary Diplomas and Certificates**ITP Summer Camp 2013****3.1.4 Licenses & Certifications**

Linux Foundation Certified Systems Administrator (LFCS-1500-0198-0100)**The Linux Foundation** Issued Feb 2015 - Expired Feb 2017**Osha 10****US Department of Labor** Issued Jul 2017**3.1.5 Skills**

Unix • Linux • System Administration • Disaster Recovery • Computing • Troubleshooting • Data Center • Firewalls • Security • Cloud • GTAW • SMAW • FCAW • GMAW

4. Rethinkeverything

4.1 There is no bullet list like MY Bullet list

(Notes to myself #rethinkeverything)

Switch hands

- Move the pain
- Rewire the brain

It's your data

- Hand Copy it in Triplicate.
- If its social then scrape it and automate it. God knows they do.

It's your work

- They can't own what you learn.
- Redact and copy your notes in Triplicate.
- Create/test and share open source gists/solutions
- Make work pathways to give back to the community
- If you have to learn it you best use it at home (lxd5/ansible/jellyfish/usw)

Work on one less thing (simplify)

- Every convenience is a point of failure or an attack surface.
- Git does not need to look good to be usefull. (--gitea, --gitlab, ++bare-git+hooks/mirrors)
- -- Twitter
- with or without musk
- also #fuckthatguy.
- If your content is usefull it will recieve the appropriate tweets, links, usw.
- and if it doesnt the internet is fundamentally broken.
- If your wysywig is so unusable you don't "blog" Throw it away. (--wordpress)

Home is where the heart is

- Don't let pi/routers do server/container work.
- PiHole (filtering dns)
- dhcp
- look at virtualized routing
- If you can't netboot off of it is it really your network.
- Same goes for centralized management (ldap).
- dhcp
- tftp
- iscsi
- All active work behind at least one firewall.

- Automate pushes (including this site).
- Streamline/cleanup html generation
- It should also be replicated in at least one other location

Let's go to your place

- Ticketing systems should be more like distributed ~~punch~~ notecards. --trac
- You shouldn't be giving out XXX bucks a month to post your public images so they can be "distributed"
- flicker free
- multi homed

Just because you have source control doesn't mean it's all code

- use markdown/git for most things.
- but focus on the english.

Work imitates life

(What problems are we trying to solve)

- Minimize technical debt both past and future.
- Disentangle the various interconnected pieces and dependencies
- Automate as much as possible
- Document what is to be done. (Specification and sample implementation)
- Practice experience based stepwise refinement.

Start making sense.

CONSOLIDATE HOME NETWORK USING OPNSENSE.

- REMOVE openwrt router
- REMOVE dedicated caching server
- REMOVE dedicated dnsmasq server.
- REMOVE (that fucking qwest router)
- KEEP Pihole-FTL dns based blacklisting
- ADD Better firewall rules
- ADD VPN access to home network
- ADD Isolated Wireless network for solar array controller.

4.1.1 Sarcasms (link them later)

1. See: Tufte's critiq of power point.
2. "as code" is only as good as managements understanding of the job of the people who actually write and maintain it minus corporate whims, the abuse of executive privilege, and cultural constraints..

4.2 Costello, an Ubuntu 22.04 lxd 5 home server.

```
# snap install lxd
# apt install htop openssh-server install netatalk zfsutils-linux
# apt remove --purge network-manager network-manager-gnome network-manager-pptp network-manager-pptp-gnome

# ip a |sed 's/^/# /'>> /etc/netplan/01-network-manager-all.yaml
# nano /etc/netplan/01-network-manager-all.yaml
#----- /etc/netplan/01-network-manager-all.yaml
#
# Dont Let NetworkManager manage *ANY* devices on this system
#
# enp3s0f0 68:fe:f7:09:3c:4c
# Dont Let NetworkManager manage *ANY* devices on this system
#2: enp3s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
#   link/ether 68:fe:f7:09:3c:4c brd ff:ff:ff:ff:ff:ff
#   inet 192.168.128.229/17 brd 192.168.255.255 scope global dynamic noprefixroute enp3s0f0
#3: wlp2s0: <NO-CARRIER,BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state DORMANT group default qlen 1000
#   link/ether 18:81:0e:ee:7c:88 brd ff:ff:ff:ff:ff:ff

network:
  version: 2
  renderer: networkd
  ethernets:
    eth0:
      match:
        macaddress: 68:fe:f7:09:3c:4c
      mtu: 7000
      dhcp4: no
      dhcp6: no
      set-name: eth0
    wlp2s0:
      dhcp4: no
      dhcp6: no
  bridges:
    br0:
      dhcp4: no
      dhcp6: no
      mtu: 7000
      addresses:
        - 192.168.129.45/17
      #gateway4: 192.168.129.1
      routes:
        - to: default
          via: 192.168.129.1
      nameservers:
        addresses:
          - 192.168.129.1
          - 198.202.31.132
      interfaces:
        - eth0
# netplan apply
# reboot
```

```
fdisk -l
# fdisk -l
...
Disk /dev/nvme0n1: 1.82 TiB, 2000398934016 bytes, 3907029168 sectors
Disk model: CT2000P2SSD8
...
Device            Start      End      Sectors  Size Type
/dev/nvme0n1p1    2048      1050623  1048576  512M EFI System
/dev/nvme0n1p2    1050624   3907028991 3905978368 1.8T Linux filesystem
...
Disk /dev/sda: 12.73 TiB, 14000519643136 bytes, 27344764928 sectors
Disk model: M001G-2KJ103
...
Device            Start      End      Sectors  Size Type
/dev/sda1          2048    6442452991 6442450944 3T Linux filesystem
/dev/sda2    6442452992 12884903935 6442450944 3T Linux filesystem
/dev/sda3    12884903936 21474838527 8589934592 4T Linux filesystem
/dev/sda4    21474838528 27344764894 5869926367 2.7T Linux filesystem
...
Disk /dev/sdb: 1.86 TiB, 2048408248320 bytes, 4000797360 sectors
Disk model: JAJ5600M2TB
...
Device            Start      End      Sectors  Size Type
/dev/sdb1          40     409639     409600  200M EFI System
/dev/sdb2    409640 4000797319 4000387680 1.9T Apple APFS
# ls -lsa /dev/disk/by-id |grep sda
# zpool create tank wmn-0x5000c500dc29d6c5-part4
```

```
lxd init
# lxd init
Would you like to use LXD clustering? (yes/no) [default=no]: yes
What IP address or DNS name should be used to reach this node? [default=192.168.129.45]:
Are you joining an existing cluster? (yes/no) [default=no]:
```

```

What name should be used to identify this node in the cluster? [default=costello]:
Setup password authentication on the cluster? (yes/no) [default=no]: yes
Trust password for new clients:
Again:
Do you want to configure a new local storage pool? (yes/no) [default=yes]:
Name of the storage backend to use (btrfs, dir, lvm, zfs) [default=zfs]:
Create a new ZFS pool? (yes/no) [default=yes]:
Would you like to use an existing empty block device (e.g. a disk or partition)? (yes/no) [default=no]: yes
Path to the existing block device: /dev/disk/by-id/wwn-0x5000c500dc29d6c5-part1
Do you want to configure a new remote storage pool? (yes/no) [default=no]:
Would you like to connect to a MAAS server? (yes/no) [default=no]:
Would you like to configure LXD to use an existing bridge or host interface? (yes/no) [default=no]: yes
Name of the existing bridge or host interface: br0
Would you like stale cached images to be updated automatically? (yes/no) [default=yes]:
Would you like a YAML "lxd init" preseed to be printed? (yes/no) [default=no]: yes
config:
  core.https_address: 192.168.129.45:8443
  core.trust_password: ON.TACOCAT.NO
networks: []
storage_pools:
- config:
  source: /dev/disk/by-id/wwn-0x5000c500dc29d6c5-part1
  description: ""
  name: local
  driver: zfs
profiles:
- config: {}
  description: ""
  devices:
    eth0:
      name: eth0
      nictype: bridged
      parent: br0
      type: nic
    root:
      path: /
      pool: local
      type: disk
  name: default
projects: []
cluster:
  server_name: costello
  enabled: true
  member_config: []
  cluster_address: ""
  cluster_certificate: ""
  server_address: ""
  cluster_password: ""
  cluster_certificate_path: ""
  cluster_token: ""

```

```

nano /etc/systemd/resolved.conf
[Resolve]
DNS=192.168.129.250
#FallbackDNS=
Domains=lan suspetdevices.com local
ln -sf /run/systemd/resolve/resolv.conf /etc/resolv.conf
nano /etc/nsswitch.conf
...
hosts:          files mdns4_minimal dns [NOTFOUND=return] dns
...

```

```

# nano /etc/netatalk/afp.conf
;----- /etc/netatalk/afp.conf
; Netatalk 3.x configuration file
;

[Global]
; Global server settings

; pretty sure this one stays.
map acls = mode
; Not sure about the next two
aclinherit = passthrough
aclmode = passthrough

[tank]
path = /tank
ea=none
# service netatalk restart

```

```

# chown feurig /tank/
# su - feurig

```

Documentation.

```
ip3 install mkdocs
pip3 install mkdocs-bootswatch
pip3 install mkdocs-multirepo-plugin
pip3 install mkdocs-mermaid2-plugin
pip3 install autolink-references-mkdocs-plugin
#mkdocs serve

# Install mkdocs
# Themes
# Multi-repo support
# Mermaid.js support
# Autolink tickets inserted into docs
```

4.3 otto (OTTO) a ubuntu laptop/server.

Recently I built out a ubuntu laptop to work on these rariton PDUs we have a bunch of at work. So I had the install media handy when I locked myself out of Shirleys retired (thanks again apple) laptop which was running octoprint. The usb network adapter had been on one of a pile of rpi-zeros with odd names like ottootto. Since the original octoprint server was named something else I was trying to figure out how the heck I had a persistan system with that name. On the other hand I love that song.

```

08:21 <+stgraber> feurig: I believe I made a video about it an MAAS some time ago
08:21 <+stgraber> feurig: anyway, it's basically:
08:22 <+stgraber> lxc init my-pxe --empty --vm
08:22 <+stgraber> lxc config device override my-pxe eth0 boot.priority=10
08:22 <+stgraber> lxc start my-pxe --console=vga
08:22 <+stgraber> the boot.priority step is to have QEMU prefer network boot over local disk
08:23 <+stgraber> you may also want to grow the root disk depending on your needs: lxc config device override my-pxe root size=50GiB

```

4.4 pure config is out of scope of this note

4.5 setting up pure -> freebsd

4.6 /etc/rc.conf.local

```
iscsid_enable="YES" iscsictl_enable="YES" iscsictl_flags="-Aa"
```

4.7 grab rc.d file from zfs-backup2:/etc/rc.d/zpool_iscsi

4.8 so system will try to import zpool after iscsi has settled

```
zpool_iscsi_enable="YES"
```

4.9 /etc/iscsi.conf

```

pure01-ct0-1 { TargetAddress = pure01-ct0-1.evergreen.laika.com SessionType = Discovery InitiatorName = iqn.
2005-06.com.laika:freebsd-hostname.evergreen.laika.com } pure01-ct0-2 { TargetAddress = pure01-ct0-2.evergreen.laika.com
SessionType = Discovery InitiatorName = iqn.2005-06.com.laika:freebsd-hostname.evergreen.laika.com } pure01-ct1-1 {
TargetAddress = pure01-ct1-1.evergreen.laika.com SessionType = Discovery InitiatorName = iqn.2005-06.com.laika:freebsd-
hostname.evergreen.laika.com } pure01-ct1-2 { TargetAddress = pure01-ct1-2.evergreen.laika.com SessionType = Discovery
InitiatorName = iqn.2005-06.com.laika:freebsd-hostname.evergreen.laika.com }

```

```
service iscsid start service iscsictl start
```

4.10 view iscsi luns

```
iscsictl -L
```

4.11 remove luns

```
iscsictl -Ra
```

4.12 add luns

```
icsictl -Aa
```

4.13 freebsd initiator doesn't handle multipath.

4.14 The geom_multipath kernel module does

4.15 create multipath device

```
kldload geom_multipath
```

4.16 make it survive a reboot

```
echo geom_multipath_load="YES" >> /boot/loader.conf
```

4.17

```
gmultipath label mp0 da4 da5 da6 da7
```

4.18 now you can create a zpool using the mp0 device

```
zpool create zjail multipath/mp0 zfs set mountpoint=/jails zjail zpool set autotrim=on zroot zfs set compression=off zjail
```


4.19 Utah

4.20 LXD5

4.20.1 Managing lxd differently

```
Hoffa:~ don$ brew install lxc
Hoffa:~ don$ lxc remote add costello.local
Certificate fingerprint: c7b38e549c397aa9d5e63489bf9a5f3987ec8d67dada692cafdecca924d4b8bf
ok (y/n/[fingerprint])? y
Admin password for costello.local:
Client certificate now trusted by server: costello.local
Hoffa:~ don$ lxc remote set-default costello.local
Hoffa:~ don$ lxc remote list
```

NAME	URL	PROTOCOL	AUTH TYPE	PUBLIC	STATIC	GLOBAL
costello.local (current)	https://costello.local:8443	lxd	tls	NO	NO	NO
images	https://images.linuxcontainers.org	simplestreams	none	YES	NO	NO
local	unix://	lxd	file access	NO	YES	NO
ubuntu	https://cloud-images.ubuntu.com/releases	simplestreams	none	YES	YES	NO
ubuntu-daily	https://cloud-images.ubuntu.com/daily	simplestreams	none	YES	YES	NO

4.21 Sense

4.21.1 Start making sense.

At work we use pf on freebsd-13.1 for our firewalls. I have been re-learning it as my freebsd-firewall experience is over 2 decades old. At home and in the colo we have been using openwrt which is great but a real pain to keep updated and deploy. I have been looking at pfsense and in the process, I discovered opnsense. If my BSD/PF chops ever get good enough I may go to straight freebsd but the convenience of guided configuration of a secure system is hard to ignore [\[1\]](#)

The hardware

As I was considering looking at pfsense I scored a pair of routers with 6x1G ports, and room for a pair of ssds.



First thing I did was to pull the os disk and replace it with a 1T ssd and upgrade the memory. The second thing I did was to replace the fans with quieter ones and print a pair of noise reducing mufflers. (I should blog about this on suspect devices at some point)

After that I installed opnsense and started working on my list of things to do.

THE GOAL: CONSOLIDATE HOME NETWORK USING OPNSENSE.

- ☒ REMOVE openwrt router
- ☒ REMOVE dedicated caching server (Done)
- ☒ REMOVE dedicated dnsmasq server
- ☐ REMOVE (that f*cking centurylink router)
- ☒ KEEP Pihole-FTL dns based blacklisting
- ☐ ADD Better firewall rules
- ☐ ADD VPN acces to home network
- ☐ ADD Isolated Wireless network for solar array controller.

FOOTNOTES/SARCASMS

[1\)](#) On the other hand having a gui make things easier makes it easy to break things and less easy to debug them. (having managed to brick the home network trying to add an isolated wireless network)

4.21.2 Centurylink fiber

Linkdump

- <https://gist.github.com/matracey/12cc7c51297561f49b4d1a95b68abc45>
- <https://forum.netgate.com/topic/83139/pppoe-on-wan-link-for-centurylink-gigabit-service/23>
- <https://www.centurylink.com/home/help/internet/modems-and-routers/third-party-modem-support-and-settings.html>
- <https://www.tp-link.com/us/support/faq/2709/>

4.21.3 Keeping Pihole-ftl while moving to opnsense.

Not sure this is the best way.

Linkpile

- <https://discourse.pi-hole.net/t/opnsense-pihole/54818>
- <https://pi-hole.net/blog/2021/09/30/pi-hole-and-opnsense/>
- <https://discourse.pi-hole.net/t/first-timer-using-opnsense-and-pi-hole-guide/61694>
- <https://github.com/pi-hole/FTL>
- TODO: look at Adguard Home >>>https://www.reddit.com/r/OPNsenseFirewall/comments/tqzjy/want_to_have_a_pihole_plugin_for_opnsense_express/

4.21.4 Initial impression.

This is mostly a note about freebsd audit and why I went with opnsense. One of my coworkers didnt like some of the coding last time he looked at opnsense, but I am willing to ignore this while I work on being able to do most of this stuff by hand.

pkg audits and updates.

Out of the box pfsense-ce (2.6..) had over 20 vulnerabilities most of them in the core parts of the system. With an older version of freebsd and no real upgrade path I thought "well obscene me, this obscenes". This was really a deal breaker.

OPNsense on the other hand came out of the box with around a dozen which after a pgk update && pkg upgrade dropped down to one. This is recent, not critical and consistent with the upgrades I have been doing at work. Bodes well.

```
root@OPNsense:~ # pkg audit -F
vulnxml file up-to-date
py39-setuptools-63.1.0 is vulnerable:
  py39-setuptools -- denial of service vulnerability
  CVE: CVE-2022-40897
  WW: https://vuxml.FreeBSD.org/freebsd/1b38aec4-4149-4c7d-851c-3c4de3a1fbd0.html

1 problem(s) in 1 installed package(s) found.
root@OPNsense:~ # freebsd-version
13.1-RELEASE-p5
```

Link pile.

- <https://forum.opnsense.org/index.php?topic=18274.0>
- <https://connortumbleson.com/2022/06/06/opnsense-wireguard-pihole/>
- <https://homegrowntechie.com/discovering-migrating-to-opnsense/>

5. Serverdocs

5.1 Systems Documentation

5.1.1 Notes, and Things to be done.

[Operations Guide for current systems](#)

Server Modernization Phase I

- Moving all legacy system functions onto separate linux containers isolated from each other.
- Use mirrored disk systems to insure that disk corruption does not lead to data corruption.
- Start giving a shit about the systems, code, and sites on them.
- Own your code/data. (If your free code hosting system is shutdown or taken over by Microsoft is it really free)
- Clean up the cruft (If it doesn't bring you joy DTMFA)

Server Modernization Phase II

- Integrate Ansible into system maintenance tasks
- Reevaluate Centos and other RPM based containers built using playbooks vs profiles/scripts/cloud-init *while maintaining current security model*
- Develop off site backup strategy.

SMP III *Make Shit Happen / Own Your Shit*

- Work on secure and efficient traffic in and out of home lans (Privoxy,DNS based ad blocking,squid etc)
- Continue to refine server operation/maintanance.
- Build Gitlab and other alternatives to trac/git and evaluate workflows.
- Deploy off site backup strategy.
- Build out content.
- Start new projects.
- Distribute data and backups over the network to home servers.
- [Document home server/network setup](#)

5.2 Portland

Most of the active work here is under (start making) [sense](#).

- [wiki:Annie Annie (File server)]
- [wiki:Mullein Mullein (Firewall/VPN server)]
- [wiki:Nigel Nigel (IOT gateway)]
- [wiki:Esp8266 Exploring the Esp8266]
- [wiki:DiskRecovery Dataloss and attempt to recover my laptops disk]

5.3 Annie

Annie is the home file server for the lan in Portland.

[[Image(wiki:Annie:Home Network Diagram.jpg, width=70%)]]

Annie's primary function is to serve the house with 7+TB of redundant disk. Her second function is to provide LXD based services.

To document.... ZFS setup.

User mapped container for serving appletalk3. [[NotesOnAppleTalk3vsUbuntu | Notes on using container to provide next release capabilities]]

5.4 Ansible Scripts

To document here. * updating containers using ansible * creating containers using ansible * backing up containers.

To add here.

- Check for and stop duplicate running containers.
- Make nightly duplicates and cleanup any duplicates created.
- (related) Shift containers between machines.
- Adapt or rewrite scripts to shift Host from kb2018 to bs2020
- Document usage of existing scripts.

5.5 BS2020 (RE)Install

NotesInstalling devstack on server left entirely too much shit everywhere. Realized that devstack should be installed in a container or vm. This page documents the reinstallation of bs2020 using the remote console and admin network.

5.5.1 Firewall Setup

Allowing access to the server is discussed in the [wiki:OpenWRT OpenWRT notes] section.

5.5.2 Loading a new os via the idrac 6

- log into idrac by browsing (<https://vpn.suspectdevices.com>)
- open the virtual console. (accept all responsibility for allowing it to run)
- launch virtual media
- attach ubuntu 16.04 server iso (on your local workstation)
- boot the iso and install the server according to either the official server install instructions or your favorite i.e. <https://ittutorials.net/linux/ubuntu/install-ubuntu-16-04-lts/>
- While booting adjust the bios settings to skip PXE booting and memory testing which takes for ever
- Let the vpn on the admin lan provide the address and network settings on the first interface (will fix later)
- Select ssh server (dns,lamp, and mail will be handled by containers anything else will be faster over the net)
- ssh into box once the os is installed.

5.5.3 Post install configuration

Make primary interface static (on admin lan)

```
root@bs2020:~# nano /etc/network/interfaces
...
# The primary network interface
auto enol
iface enol inet static
    address 192.168.1.158/24
    gateway 192.168.1.1
    dns-nameservers 192.168.1.1 198.202.31.132 198.202.31.141
    dns-search vpn suspectdevices.com digithink.com
...
root@bs2020:~#
```

Update server

```
feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# apt-get update
... Done
root@bs2020:~# apt-get dist-upgrade
root@bs2020:~# apt-get install openssl-server
```

Add second admin user

```
root@bs2020:~# useradd -m joe -c"Joe Dumoulin" -Gsudo,root
root@bs2020:~# su - joe
joe@bs2020:~$ nano
joe@bs2020:~$ mkdir .ssh
joe@bs2020:~$ nano .ssh/authorized_keys
```

paste key from vpn /etc/dropbear/authorized_keys

Set initial password so that admin can sudo.

```
root@bs2020:~# vipw -s
... paste hash from medea ...
```

Consider removing password based ssh authentication once both admins can connect.

5.5.4 LXC

_ This should probably move to its own section once stable _

We want to do 3 things with lxc. * create a public facing server for dns/email/and other services which is isolated from other containers and can not access the host directly * create a similarly isolated server for openstack/devstack that can be uninstalled and which will not shit all over everything. (Attempting to containerize devstack was as disastrous as trying to uninstall it) * create user space containers for experimentation which are in themselves isolated from everything else.

LXC and the first infrastructure container

Lxd is installed but lxc is not. Install lxc lxc templates bridge utilities and zfs. In the example below we leverage lxd to create the zfs pool and to point the lxc network to the existing bridge. Once we work enough with LXC/LXD and zfs to identify the relative merits of each approach I will backfill how to do these tasks manually.

```
root@bs2020:~# sudo apt-get install lxc lxc-templates wget \
    zfsutils-linux bridge-utils ebttables openvswitch-common
...
root@bs2020:~# nano /etc/network/interfaces

# The primary network interface
auto eno1
iface eno1 inet static
    address 192.168.1.158/24
    gateway 192.168.1.1
    dns-nameservers 192.168.1.1 198.202.31.132 198.202.31.141
    dns-search vpn.suspectdevices.com digithink.com

auto br0
iface br0 inet static
    address 0.0.0.0
    bridge_ports eno4

iface eno4 inet manual

...
root@bs2020:~# lxd init
Name of the storage backend to use (dir or zfs) [default=zfs]:
Create a new ZFS pool (yes/no) [default=yes]? yes
Name of the new ZFS pool [default=lxd]: lxd4infra
Would you like to use an existing block device (yes/no) [default=no]? yes
Path to the existing block device: /dev/sde1
Would you like LXD to be available over the network (yes/no) [default=no]?
Do you want to configure the LXD bridge (yes/no) [default=yes]? no
....
root@bs2020:~# dpkg-reconfigure -p medium lxd
... no yes br0...
Warning: Stopping lxd.service, but it can still be activated by:
    lxd.socket

root@bs2020:~# lxc-create -n naomi -t ubuntu -B zfs --zfsroot=lxd4infra
lxc.rootfs = /var/lib/lxc/naomi/rootfs
lxc.rootfs.backend = zfs
lxc.utsname = naomi
lxc.arch = amd64
..
root@bs2020:~# nano /var/lib/lxc/naomi/config
..... check network ....
# Network configuration
lxc.network.type = veth
lxc.network.link = br0
lxc.network.flags = up
lxc.network.hwaddr = 00:16:3e:dc:6d:b4
# Assign static IP Address (currently done by container)
#lxc.network.ipv4 = 192.168.1.161/24
#lxc.network.ipv4.gateway = 192.168.1.1
..... add this ....
# Autostart
lxc.start.auto = 1
lxc.start.delay = 5
lxc.start.order = 100
....
root@bs2020:~# reboot
```

adding admin users and basic services (lock ubuntu user before starting network)

```
root@bs2020:~# lxc-attach -n naomi
root@naomi:~# passwd -l ubuntu
```

```

root@naomi:~# vi /etc/network/interfaces
... add the following ...
auto eth0
iface eth0 inet static
    address 198.202.31.142/25
    gateway 198.202.31.129
    dns-nameservers 198.202.31.132 198.202.31.141 8.8.8.8
    dns-search vpn suspectdevices.com digithink.com

root@naomi:~# ifdown eth0 && ifup eth0
root@naomi:~# ping digithink.com
root@naomi:~# apt-get update
root@naomi:~# apt-get install openssl-server nano
root@naomi:~# useradd -Gsudo,root -m -c"Donald Delmar Davis" feurig
root@naomi:~# useradd -Gsudo,root -m -c"Joe Dumoulin" joe
root@naomi:~# vipw -s
... paste hash from other system...
root@naomi:~# tail -2 /etc/passwd >passwd.add
root@naomi:~# tail -2 /etc/shadow >shadow.add
root@naomi:~# tar -czvf fnj.tgz /home
root@naomi:~# exit
root@bs2020~# cp /var/lib/lxc/naomi/rootfs/root/*.add ~feurig/
root@bs2020~# cp /var/lib/lxc/naomi/rootfs/root/fnj.tgz ~feurig/

```

tuning bs2020

TODO: <https://github.com/lxc/lxd/blob/master/doc/production-setup.md>

5.5.5 devstack lxc container (FAIL)

This does not work. As far as I can tell you can only install devstack on raw hardware and let it install all of its ever moving dependencies. I was able to do this where pike was 6 months ago but not uninstall and reinstall is using the same version.

I don't believe it can trust anything this moving to be sane let alone secure.

SEE: GoodByeOpenstack

I may attempt this again within a KVM once I establish that the KVM framework is securable and that it will play nice with the existing containers.

5.5.6 LXD Container and Docker Install

SEE: [wiki:LXDContainerWithDockerNotes Creating LXD Container with static ip and Docker Profile]

lxc docker references

- <https://www.flockport.com/lxc-vs-docker/>
- <https://www.upguard.com/articles/docker-vs-lxc>
- <http://www.zdnet.com/article/ubuntu-lxd-not-a-docker-replacement-a-docker-enhancement/>
- <https://stackoverflow.com/questions/37227349/unable-to-start-docker-service-in-ubuntu-16-04>
- <https://stackoverflow.com/questions/32002882/error-starting-docker-daemon-on-ubuntu-14-04-devices-cgroup-isnt-mounted>

control groups / other related references

- <https://help.ubuntu.com/lts/serverguide/cgroups-overview.html>
- <https://askubuntu.com/questions/836469/install-cgconfig-in-ubuntu-16-04>
- <https://help.ubuntu.com/lts/serverguide/cgroups.html>

LXC REFERENCES

- <https://www.ubuntu.com/containers/lxd>
- <https://insights.ubuntu.com/2016/04/07/lxd-networking-lxdbr0-explained/>
- <https://bayton.org/docs/linux/lxd/lxd-zfs-and-bridged-networking-on-ubuntu-16-04-lts/>
- <https://www.simpleprecision.com/ubuntu-16-04-lxd-networking-simple-bridge/>

- <https://askubuntu.com/questions/453659/lxc-containers-fail-to-autoboot-in-14-04-trusty-using-lxc-start-auto-1>
- <https://help.ubuntu.com/lts/serverguide/lxc.html>
- <http://www.itzgeek.com/how-to/linux/ubuntu-how-to/setup-linux-container-with-lxc-on-ubuntu-16-04-14-04.html>
- <https://bayton.org/docs/linux/lxd/lxd-zfs-and-bridged-networking-on-ubuntu-16-04-lts/>
- <https://stgraber.org/2016/03/15/lxd-2-0-installing-and-configuring-lxd-212/>
- <https://wiki.ubuntu.com/LxcSecurity>
- <https://insights.ubuntu.com/2016/03/16/lxd-2-0-installing-and-configuring-lxd-212/>

5.5.7 fuckups

- openstack/devstack shits all over your server you uninstall it by starting over
- CHECK TO MAKE SURE YOU ARE IN A CONTAINER BEFORE INSTALLING THE POS THE BARE METAL INTALL IS TOLERABLE BUT NOT FUN.
- installing the virtual server host installs KVM and its kernel. uninstalling it leaves you with a kernel that can't find the network.
- don't press f10 during boot whatever you do and if you do follow this... <http://crtech.tips/lifecycle-controller-hanging-during-post/>
- do not give br0 an address as it will then become a public facing interface with direct access to the host server.
- local.conf password can't contain any shell characters (%\$@!) much like the puppet installer..
- host must also have bridge tables (ebtables) and openvswitch installed.
- kernel modules needed in lxc containers need to be installed in the host.
- deleting container zfs pool and storage without telling lxd not to use it is problematic. Hint

```
root@bs2020:~# lxc config show config: storage.zfs_pool_name: lxd4dev
```

5.6 Bleeding Edge (old)

Looking at LTS debian/ubuntu 18.04 for the next 5 years (bs2020/phillip)

As an excersize we ran up a 17.10 ubuntu container to see what all was going to break when we upgraded. So far the usual suspects (Network configuration, startup etc) are all fucked up. We upgraded this to 18.04 using do-release-upgrade.

```
root@phillip:~# do-release-upgrade -d
```

Then we started working on the bullshit.

Stupid Idea #1 netplan

According to the Release Notes for Bionic Beaver: on top of adding color emojis they rewrote the network management layer based on the worst innovations of modern linux (systemd and NetworManager)

"Netplan is a YAML network configuration abstraction for various backends (NetworkManager, networkd).

It is a utility for easily configuring networking on a system. It can be used by writing a YAML description of the required network interfaces with what they should be configured to do. From this description it will generate the required configuration for a chosen renderer tool.

Netplan reads network configuration from /etc/netplan/*.yaml which are written by administrators, installers, cloud image instantiations, or other OS deployments. During early boot it then generates backend specific configuration files in /run to hand off control of devices to a particular networking daemon."

In otherwords we are ripping up everything and hoping the the details will work themselves out even though they are not defined and buried in several layers of bullshit written by children and adult children (zb your average modern CTO).

5.6.1 Making it work

After a lot of digging I edited this file on phillip and rebooted the container.

```
root@phillip:~# nano /etc/netplan/50-cloud-init.yaml
# This file is generated from information provided by
# the datasource.  Changes to it will not persist across an instance.
# To disable cloud-init's network configuration capabilities, write a file
# /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg with the following:
# network: {config: disabled}
network:
  version: 2
  ethernets:
    eth0:
      dhcp4: no
      addresses: [198.202.31.223/25]
      gateway4: 198.202.31.129
      nameservers:
        search: [suspectdevices.com fromhell.com vpn]
        addresses: [198.202.31.141]
```

Note that even this file says its generated rather than referenced. Really? 'provided by "the datasource"' WHAT DATASOURCE???? FUCKING KIDS. The datasource in this case would be cloud init. In this case cloud init has been told not to configure the network.

5.6.2 linkdump

- <https://wiki.ubuntu.com/Netplan/Design>
- <https://wiki.edubuntu.org/BionicBeaver/ReleaseNotes>
- <http://www.ubuntugeek.com/how-to-assign-static-ip-address-in-ubuntu-17-10-artful-aardvark.html>
-

5.7 CloudServerConfiguration

Videoranch Cloud Server Configuration.

The purpose of this document is provide information on how gihon.orgs cloud server is currently configured and basic guidelines for maintaining it.

# Date	# Author	# Email	# Comments
28MAY16	Donald Delmar Davis	don@suspectdevices.com	Initial document

Background

We were asked to convert a 15 year old internet server running freebsd to the cloud. We started by setting up a staging server running Ubuntu 14.04 and migrating the users data and log files from the old server. This provided a backup of the original data and a place where we could work without having to pay for disk or bandwidth before deploying the final product. After a long process of porting all of the users and web sites that the server had served over the decades we began identifying which services, users, and domains were needed on the server. Given a much smaller set of users and web sites that were actually needed, we deployed an AWS image based on the AMI provided by the commercial entity which maintains Ubuntu. The active users users and web content have been installed on this server and the remainder has been archived to an external disk.

5.7.1 The Base Image

We chose to deploy an image provided by Canonical specifically for AWS "ubuntu-trusty-14.04-amd64-server-20150325 (ami-5189a661)" <http://cloud-images.ubuntu.com/releases/trusty/release-20150325/>

Adjustments to the image

The ubuntu user which provides a back door through which AWS allows users that it has authenticated to have root access to the instance. Unfortunately the ubuntu UID(1000) was already taken (jess) so it was moved to 999 and files owned by it were migrated as well.

```
chown --from=1000:1000 999:999 / . -Rv
```

Also the mail spool was somewhere new (/var/spool/mail) so I linked the new location back to /var/mail

Additions to the image

a lamp stack was added to the image using the "tasksel" package which bundles most services into supported configurations and deploys them along with all of their dependencies. (Note that the Ubuntu Cloud Image was already installed)

```
# tasksel
Package configuration
```

```

┌──────────┴──────────┐ Software selection ┌──────────┴──────────┐
| You can choose to install one or more of the following predefined collections of software. |
|
```

```
| Choose software to install:
```

```
|
|  [*] Basic Ubuntu server
|  [*] OpenSSH server
|  [ ] DNS server
|  [*] LAMP server
|  [*] Mail server
|  [*] PostgreSQL database
|  [ ] Print server
|  [ ] Samba file server
|  [ ] Tomcat Java server
|  [*] Ubuntu Cloud Image (instance)
|  [ ] Virtual Machine host
|
```

```
...
```

```
<Ok>
```


users and superusers

The following users were added to the system.

```
jess:x:1000:1000:Jessica Kent:/home/jess:/bin/csh
gepr:x:1053:1053:Glen E Ropella:/home/gepr:/bin/bash
don:x:1054:1054:Donald Delmar Davis:/home/don:/bin/bash
vic:x:1002:1002:Victoria Kennedy:/home/vic:/bin/bash
nez:x:1003:1003:Michael Nesmith:/home/nez:/bin/bash
vranch:x:1004:1004:Videoranch User:/home/vranch:/bin/bash
foreman:x:1005:1005:Videoranch Foreman:/home/foreman:/bin/tcsh
navajoslim:x:1007:1007:Navajo Slim:/home/navajoslim:/bin/bash
gihon:x:1017:1017:Gihon Foundation:/home/gihon:/bin/bash
vk:x:1021:1021:Victoria Kennedy:/home/vk:/bin/bash
vrresume:x:1024:1024:videoranch resume:/home/vrresume:/bin/bash
vak:x:1027:1027:victoria kennedy:/home/vak:/bin/tcsh
nezrays:x:1031:1031:nezrays:/usr/home/vranch/nezrays/www:/bin/sh
vr3d:x:1035:1035:VR3D:/home/vr3d:/bin/sh
staging:x:1041:1041:staging:/home/staging:/bin/bash
nesmith:x:1042:1042:nesmith:/home/nesmith:/bin/bash
director:x:1045:1045:Jessica Kent:/home/director:/bin/bash
petetest:x:1048:1048:petetest:/home/petetest:/bin/bash
mn:x:1022:1022:Michael Nesmith:/home/mn:/bin/bash
```

This had to be done manually as some of the original passwords were so old that their encryption methods were no longer supported. In cases where the users were less than a few years old the users passwords transferred to the new system seamlessly. In other cases the passwords will have to be reset by someone with root access.

```
ubuntu@cloud # passwd vranch
```

Their mail spools (/var/mail/), and home directories were copied over as well.

sudo privileges were enabled for members of the sudo group.

```
ubuntu@cloud # vigr
...
sudo:x:27:ubuntu,jess,foreman,don,gepr
...
```

5.7.2 Apache Configuration

In addition to the home directories of the remaining users the /home/vranch directory tree and /home/gihon were copied to the new server. The server configurations were ported to be as close to the originals as possible. (exceptions noted below)

The default server is set to www.gihon.com and is configured based on the original virtual-host. The php information and much about the apache server can be queried directly at <http://videoranch.com/test.php>

```
#ServerName www.gihon.com
<VirtualHost *:80>
    ServerName www.gihon.com
    ServerAlias gihon.com www.gihon.org gihon.org cloud.gihon.com
    ServerAdmin info@digitaloffspring.com
    DocumentRoot /home/gihon/www
    <Directory '/home/gihon'>
        AllowOverride All
    </Directory>
    ScriptAlias /cgi-bin/ /home/gihon/cgi-bin
    CustomLog /home/gihon/logs/gihon-access_log common
    ErrorLog /home/gihon/logs/gihon-error_log
</VirtualHost>
```

- Note that the log files are left in user space (off of /home) this allows clients to pull and view the log files in the same way that they update the content of their web site (ftp etc)
- Some configuration directives are no longer supported and are commented out.
- Extremely dangerous statements such as AllowOverides for the root directory were modified.

All other servers are named virtualhosts. The first of which is www.videoranch.com defined in /etc/apache2/sites-enabled/www.videoranch.com.conf

```
<VirtualHost *:80>
    ServerName www.videoranch.com
    ServerAlias videoranch.com www.videoranch.com
#
    Header append p3p 'CP="OTI DSP COR CUR UNI" polyref="/w3c/p3policy.xml"'
    ServerAdmin info@digitaloffspring.com
```

```

DocumentRoot /home/vranch/videoranch/www
ScriptAlias /cgi-bin/ /home/vranch/videoranch3d/cgi-bin/
ErrorLog /home/vranch/logs/www.videoranch.com-error_log
CustomLog /home/vranch/logs/www.videoranch.com-access_log common
<Directory /home/vranch/videoranch/www>
    Options Indexes FollowSymLinks
    AllowOverride All
</Directory>
</VirtualHost>

```

5.7.3 Pro-ftp Configuration

We configured proftpd (which we vetted as a viable and secure ftp daemon) as closely as possible to the original configuration on the old server. Because AWS instances are in their own private network and access has to be explicitly allowed you must specify the PASV ports in `/etc/proftpd/proftpd.conf`. These ports must be opened up in the "Security Group" configuration as well.

```

# In some cases you have to specify passive ports range to by-pass
# firewall limitations. Ephemeral ports can be used for that, but
# feel free to use a more narrow range.
PassivePorts 49152 49153

```

Ftp in its native form is insecure and so we would prefer to have configured an SSL certificate and require TLS for all ftp requests. We were able to verify that SFTP (ftp provided by ssh).

5.7.4 Network and "Security Group" configuration

The AWS instance is placed in a private network. This network provides the instance a private ip through dhcp. For this reason the main interface is configured as follows in `/etc/networks/interfaces.d/eth0`

```

# The primary network interface
auto eth0
iface eth0 inet dhcp

```

This address is attached to the outside world via an "Elastic" ip (52.34.143.142). To connect the external traffic to the private address you have to create a "Security group" and define the rules which allow traffic in and out of the private network.

INBOUND RULES

protocol	# family	# port	# allow from
HTTP	TCP	80	0.0.0.0/0
SSH	TCP	22	0.0.0.0/0
SMTP	TCP	25	0.0.0.0/0
Custom TCP Rule	TCP	20 - 21	0.0.0.0/0
IMAP	TCP	143	0.0.0.0/0
Custom TCP Rule	TCP	49152 - 49153	0.0.0.0/0
HTTPS	TCP	443	0.0.0.0/0

Outbound rules allow all outgoing traffic.

5.7.5 Unused Capabilities

MySQL and PostgreSQL

While the M in LAMP is MySQL, Many developers prefer Postgres which is much more standards oriented and robust. Both databases are available and PHP is configured for them. At one point mysql was on the old server however neither gihon nor the model files served by videoranch.com seemed to use it. _ Note that if either database is used a mechanism to back up the data must also be implimented_

Postfix and Dovecot

The standard SMTP (email) server for most current operating systems is Postfix. The Mail server task also includes Dovecot which provides both POP and IMAP servers for clients to download any mail still on the server. To use the pop server will require the addition of the ports for pop (110) to be added to the security group configuration. _These servers are not currently configured. _

5.7.6 Log Rotation Configuration

On the previous server most log files were larger than the content being provided. Ubuntu provides a log rotation utility designed to compress and delete logs in a reasonable manner preventing them from consuming system resources over time. Since the apache logs on this system are in "user space" and not under /var/log/apache2 their location needed to be configured.

Here is the section added to /etc/logrotate.d/apache2 for the gihon.com

```
/home/gihon/logs/*_log {
    weekly
    missingok
    rotate 52
    compress
    delaycompress
    notifempty
    create 640 root adm
    sharedscripts
    postrotate
        if /etc/init.d/apache2 status > /dev/null ; then \
            /etc/init.d/apache2 reload > /dev/null; \
        fi;
    endscript
    prerotate
        if [ -d /etc/logrotate.d/httpd-prerotate ]; then \
            run-parts /etc/logrotate.d/httpd-prerotate; \
        fi; \
    endscript
}
```

5.7.7 unattended upgrades (security only)

The system is configured to automatically install security upgrades as released by the operating system. *In the event that an error occurs mail is sent to the foreman account.*

5.7.8 Operations Guide

Given the state of the previous system the soundest approach is to automate as much of the systems upkeep as possible. Log rotation and unattended system upgrades along with other minor adjustments (turning on apt's auto-remove for instance) should enable us to think of the box more as an appliance.

Backing up Server work with Live Snapshots

AWS allows a server to be backed up while running. These snapshots can be run up as separate servers (for development or to do a major release upgrade) Or they can be reattached to an existing instance (in the case of disaster or compromise). Please make a snapshot of the server whenever significant work has been done to it.

Backing up your data

Since the servers web content is in the user space. Log files, websites and other data served should be copied to a local server preferably one behind a firewall. *In particular Gihon should take care to keep updated copies of /home/gihon and /home/vranch*

Accessing the server

Privileged access can be granted through AWS to the Ubuntu user. For instructions on how to do this see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AccessingInstancesLinux.html>. The server has been configured to allow ssh access directly.

```
$ ssh www.videoranch.com
Welcome to Ubuntu 14.04.4 LTS (GNU/Linux 3.13.0-85-generic x86_64)
```

```
* Documentation: https://help.ubuntu.com/

System information as of Mon Apr 11 18:12:10 UTC 2016

System load:  0.0           Processes:      139
Usage of /:   69.8% of 29.39GB Users logged in:  1
Memory usage: 28%          IP address for eth0: 172.31.16.108
Swap usage:   0%

Graph this data and manage this system at:
https://landscape.canonical.com/

Get cloud support with Ubuntu Advantage Cloud Guest:
http://www.ubuntu.com/business/services/cloud

0 packages can be updated.
0 updates are security updates.

You have new mail.
Last login: Mon Apr 11 16:11:57 2016 from 71-34-91-188.ptld.qwest.net
don@cloud:~$
```

References

- why ubuntu? <https://insights.ubuntu.com/2014/04/15/ubuntu-14-04-its-the-cloud-platform-of-choice/>
- <https://www.digitalocean.com/community/tutorials/how-to-configure-logging-and-log-rotation-in-apache-on-an-ubuntu-vps>
- <https://help.ubuntu.com/lts/serverguide/automatic-updates.html>
- <https://anturis.com/linux-server-maintenance-checklist/>

5.8 CloudServerDocs

See [wiki:"CloudServerConfiguration"].

5.9 COBOL / Postgres / Open Enterprise/Government (Rough Draft)

I started thinking about Cobol recently. It's the only C I ever got in a Computer Science class. There are a few things that lead me to that including my experiences with Oracle VS Postgres. and how Every government from the state on down has been ripped off and bullied by Oracle and the consultants that rely on Oracle to make money. This has driven most small governments to migrate to SQLServer on Micro\$oft (King County in the early 2000s, Washington county just recently), In my opinion dangerous at best (Wanna Cry).

Recently Amazon demonstrated what most of us have known for decades. Postgresql is an enterprise capable database on par with Oracle, DBII and SQL Server. Once they made it closed source enough to monetize they used their Postgres derivative to eliminate their use of Oracle products.

While this puts Amazon in an extremely powerful position on par with Oracle, Microsoft, and IBM in relation to Enterprise and Government organizations it seems more appropriate to take the lessons they present and think seriously about business and government software on securable systems which are not owned by a Monopolistic gang of bullies who are regularly cost taxpayers Billions of dollars while not performing (Oracle VS OHP) or creating serious security concerns (Microsoft).

Open Source COBOL, Postgresql and Linux

This is what I want to explore.

CREATING LXD CONTAINERS FOR GNUCOBOL 2.2

It is my intention to create 2 containers one running Ubuntu-LTS (18.04) and the current LTS like version of Centos (7)

5.9.1 Ubuntu

GNUCobol claims that Ubuntu 18.04 will install version 2.2 in their documentation however on 18.04 only OPENCobol (1.1) is in the default repos. 2.2 is the default on Ubuntu 19.04 so we can install it using [TaskFastForwardSelectUbuntuPackages this method.]

5.9.2 Status

Currently I am having issues with any of the 3 SQL precompilers to work with GnuCobol and postgres on ubuntu (18.04) or osx (mohave). I should document this rathole.....

5.9.3 Linkdumb

- I am not alone :) <http://www.simotime.com/sys76p01.htm>

5.10 Containership Creation

- Create raid m1+0 array on machine.
- install ubuntu server + LXD (3.5) snap via ubuntu server live iso.
- fix network configuration
- add boot/serial console configuration
- add users and home directories for admin users
- install zfs
 - apt-get install nfs-kernel-server samba-common-bin zfsutils-linux
- clean drives
- initialize container data pools

5.11 Updating Hosts Notes

Updating hosts manually

The process for updating hosts is handled via apt-get in 3 steps. [#2 (2)] 1. update == check repos 2. dist-upgrade == apply updates 3. autoremove == clean up

```
root@bs2020:~# apt-get update&&apt-get dist-upgrade&& apt-get autoremove
```

Updating debian hosts using Ansible (via lxd connection)

From kb2018 we can use the apt module to update out hosts and containers however this is apt specific

```
ansible pets -m apt -a "force_apt_get=yes upgrade=yes update_cache=yes autoremove=yes"
```

Updating containers using an os agnostic script

UPDATE.SH

The ansible module is nice but specific to the operating system. To extend this to other distributions we can use the following script.

```
#!/bin/bash
# update.sh for debian/ubuntu/centos (copyleft) don@suspecdevices.com
echo ----- begin updating `uname -n` -----
if [ -x "$(command -v apt-get)" ]; then
    apt-get update
    apt-get -y dist-upgrade
    apt-get -y autoremove
fi
if [ -x "$(command -v yum)" ]; then
    echo yum upgrade.
    yum -y upgrade
fi
if [ -x "$(command -v zypper)" ]; then
    echo zypper dist-upgrade.
    zypper dist-upgrade
fi
echo =====### done=====
```

Deploying the script using either basic shell commands or shell/awk is fairly straight forward.

```
root@kb2018:~# lxc list -c n --format csv|awk '{print "lxc file push /usr/local/bin/update.sh " $1 "/usr/local/bin/"}' |bash
```

Since the hosts are on a private lan they are configured to trust each other. This means that the above deployment can be pushed to the other server as well.

```
root@kb2018:~# lxc list bs2020: -c n --format csv|awk '{print "lxc file push /usr/local/bin/update.sh bs2020:" $1 "/usr/local/bin/"}' |bash
```

Running the script using is equally simple.

```
root@kb2018:~# for h in `lxc list bs2020: -c n --format csv`; do lxc exec bs2020:$h update.sh; done
root@kb2018:~# for h in `lxc list local: -c n --format csv`; do lxc exec local:$h update.sh; done
```

Ansible improves on this simplicity using the file and raw modules.

```
.... fill in file deployment example ....
root@kb2018:/etc/ansible# ansible pets -m copy -a 'src=/etc/ansible/files/update.sh dest=/usr/local/bin/ owner=root group=root mode=0774'
root@kb2018:~# ansible pets -m raw -a "update.sh"
```

Currently only current ubuntu (18.04) is downloaded for creating hosts. Any other hosts will require manual configuration or use of a custom profile.

for example (more info at TaskCreatingNewContainers)

```

root@bs2020:~# lxc image list kb2018:
+-----+-----+-----+-----+-----+-----+-----+
| ALIAS | FINGERPRINT | PUBLIC | DESCRIPTION | ARCH | SIZE | UPLOAD DATE |
+-----+-----+-----+-----+-----+-----+-----+
| ubuntu-lts | ae465acff89b | no | ubuntu 18.04 LTS amd64 (release) (20180613) | x86_64 | 173.14MB | Jun 16, 2018 at 10:07pm (UTC) |
+-----+-----+-----+-----+-----+-----+-----+
root@bs2020:~# lxc init kb2020:ubuntu-lts test18 -p susdev19
Creating test18
root@bs2020:~# lxc start test18
root@bs2020:~# lxc exec test18 bash
root@test18:~# nano /etc/netplan/50-cloud-init.yaml
network:
  version: 2
  ethernet:
    eth0:
      dhcp4: no
      addresses: [198.202.31.216/25]
      gateway4: 198.202.31.129
      nameservers:
        search: [suspectdevices.com fromhell.com vpn]
        addresses: [198.202.31.141]
root@test18:~# reboot
root@test18:~# update.sh
root@test18:~# root@bs2020:~# lxc list
+-----+-----+-----+-----+-----+-----+-----+
| NAME | STATE | IPV4 | IPV6 | TYPE | SNAPSHOTS |
+-----+-----+-----+-----+-----+-----+-----+
....
+-----+-----+-----+-----+-----+-----+-----+
| test18 | RUNNING | 198.202.31.216 (eth0) | | PERSISTENT | 0 |
+-----+-----+-----+-----+-----+-----+-----+
root@bs2020:~#

```

5.11.1 linkdump

- <https://lxd.readthedocs.io/en/latest/backup/>
- <https://s3hh.wordpress.com/2016/05/08/using-lxd-snapshots/>
- <https://blog.ubuntu.com/2015/03/20/installing-lxd-and-the-command-line-tool>
- <https://www.virtualizationhowto.com/2018/09/installing-and-configuring-ubuntu-server-18-04-lts/>
- <http://blog.dustinkirkland.com/2018/02/rfc-new-ubuntu-1804-lts-server-installer.html>
- <https://blog.printk.io/2018/04/ubuntu-18-04-lts-bionic-beaver-server-installer-differences/>
- <https://github.com/lxc/lxd/issues/4526>
- <https://github.com/lxc/lxd/issues/4619>
- <https://docs.oracle.com/cd/E19253-01/819-5461/gbcya/index.html>
- <http://lxd.readthedocs.io/en/latest/backup/#container-backup-and-restore>
- <https://stgraber.org/2016/03/30/lxd-2-0-image-management-512/>
- <https://github.com/lxc/lxd/issues/2669>
- <https://github.com/lxc/lxd/issues/3730>
- <https://www.thegeekdiary.com/zfs-tutorials-creating-zfs-snapshot-and-clones/>
- <https://pthree.org/2012/12/19/zfs-administration-part-xii-snapshots-and-clones/>
- <https://serverfault.com/questions/74411/best-compression-for-zfs-send-receive>
- <http://everycity.co.uk/alsadair/2010/07/using-mbuffer-to-speed-up-slow-zfs-send-zfs-receive/>
- <http://www.polyomica.com/improving-transfer-speeds-for-zfs-sendreceive-in-a-local-network/>

Foot Notes

.... review / decruft .. [= #fn1 1]) The original purpose of this server was to evaluate openstack.

Openstack requires relinquishing complete control of the host server to an overtly complicated pile of layers which once installed cannot be uninstalled without completely re-installing the entire operating system. This is not that unusual (my first installation of puppet was equally badly behaved and destructive) but it does not instill confidence in software with cart blanch access to everything.

See: [wiki:GoodByeOpenstack]

Our search for a way to deploy such an insecure POS led us to look deeply into the lightweight container system provided by lxc. We attempted to create an isolated server for openstack/devstack that can be uninstalled and which will not shit all over everything. (Attempting to containerize devstack was as disastrous as trying to uninstall it)

In the process we discovered a way to create a public facing server for dns/email/and other services which is isolated from other containers and can not access the host directly.

By extending this new set of tools we are also able to create user space containers for experimentation which are in themselves isolated from everything else.

[= #fn2 2]) Ubuntu can be configured to auto update however in my experience this leads to a false sense of security and a lack of awareness of what is broken/changing. Also, when autoupdates fail they do not recover gracefully, will not apply the next set of updates, and it's a major pain in the ass to fix them. For this reason I tend to use apticron to notify us when updates are available and manually update them.

For BS2020 and naomi, I also tend to look at what is being done instead of adding the -y parameter to apt-get.

[= #fn3 3]) There are some side effects to this method for instance moving to a new server can apply the other servers default profile to it. I have also noticed that moving from a snapshot to a new container starts the new container.

5.12 Old DL380 Raid Notes

5.12.1 Problem: Where are my disks???

When we installed the os on our new (to us) proliant DL380, Only a single disk was visible in spite of there having been 6 disks installed. This is because the DL380s disk controller was set up in raid mode and did not expose disks until they were configured as "logical" disks.

This is unlike the Dell PowerEdge we have which detects and presents the drives in a hot swappable fashion while still allowing some disks to participate in raid arrays.

Since we use hardware raid mirroring on the boot disks, Adding, removing or replacing disks requires configuring the raid controller.

5.12.2 Configuring the disks using the raid controller bios

... use some words here ...

```
steve:~ don$ ssh -p 2222 feurig@vpn.suspectdevices.com
User:feurig logged-in to kb2018.suspectdevices.com(192.168.31.119 / FE80::9E8E:99FF:FE0C:BAD8)
iLO 3 Advanced for BladeSystem 1.88 at Jul 13 2016
Server Name: kb2018
Server Power: On



hpiLO-> vsp

Virtual Serial Port Active: COM2

Starting virtual serial port.
Press 'ESC (' to return to the CLI Session.

root@kb2018:~# fdisk -l|grep Disk\ \
Disk /dev/loop0: 86.9 MiB, 91099136 bytes, 177928 sectors
Disk /dev/loop1: 87.9 MiB, 92164096 bytes, 180008 sectors
Disk /dev/loop2: 63.4 MiB, 66486272 bytes, 129856 sectors
Disk /dev/sda: 136.7 GiB, 146778685440 bytes, 286677120 sectors

root@kb2018:~# reboot
[ OK ] Stopped Stop ureadahead data collection 45s after completed      Stopping Session 98 of user feurig.
      Stopping Availability of block devices...
...
[ OK ] Reached target Shutdown.
[ OK ] Reached target Final Step.
      Starting Reboot...
[292357.910620] reboot: Restarting system
```

After several seconds you will see a text based bios screen  After the network controller is started the raid controller will give you a chance to configure it. **_PRESS F8 NOW!! _** 

If you miss it you will have to escape back to the ILO3 and power cycle the machine. *(This is ok because the disks are not active until the machine actually boots)*

```
Booting from Hard Drive C:
<ESC> (
hpiLO-> power off hard

status=0
status_tag=COMMAND COMPLETED
Wed Sep 26 15:31:57 2018

Forcing server power off .....
Please wait 6 seconds for this operation to complete.

hpiLO-> power

status=0
status_tag=COMMAND COMPLETED
Wed Sep 26 15:32:04 2018
```

```
power: server power is currently: Off
```

```
hpiL0-> power on
```

```
status=0
status_tag=COMMAND COMPLETED
Wed Sep 26 15:32:21 2018
```

```
Server powering on .....
```

```
hpiL0-> vsp
```

```
Virtual Serial Port Active: COM2
```

```
Starting virtual serial port.
Press 'ESC (' to return to the CLI Session.
```

Once in the raid controller bios you will get a main menu.

[[Image(CaptiveRaidController:ViewLogicalDrive.png)]]

If you select view logical drives will see that the first two disks are combined into a mirrored pair and that there are no other drives defined.

So we select "Create Logical Drive". Which gives us the following screen.

[[Image(CaptiveRaidController:CreateLogicalDriveDefaults.png)]]

Notice that the defaults are to create a raid 1+0 array with the first two matching disks. Deselecting either disk (down arrow, spacebar) will cause the raid configuration to automatically drop to RAID 0

Press Enter when finished. The next screen will ask you to verify the creation

Repeat this for each remaining disk.

When you are finished you can view the logical drives. [[Image(CaptiveRaidController:RaidConfFinished.png)]]

The key will walk you back out so you can continue to boot.

5.12.3 success

```
root@kb2018:~# fdisk -l|grep Disk\ \ /
Disk /dev/loop0: 86.9 MiB, 91099136 bytes, 177928 sectors
Disk /dev/loop1: 87.9 MiB, 92164096 bytes, 180008 sectors
Disk /dev/loop2: 63.4 MiB, 66486272 bytes, 129856 sectors
Disk /dev/sda: 136.7 GiB, 146778685440 bytes, 286677120 sectors
Disk /dev/sdb: 223.6 GiB, 240021504000 bytes, 468792000 sectors
Disk /dev/sdc: 223.6 GiB, 240021504000 bytes, 468792000 sectors
Disk /dev/sdd: 279.4 GiB, 299966445568 bytes, 585871964 sectors
Disk /dev/sde: 279.4 GiB, 299966445568 bytes, 585871964 sectors
root@kb2018:~#
```

5.12.4 Using HP utilities to configure the controller without downing the server

HP provides utilities and officially supports bionic and hosts a repo for it. It includes a server that can be accessed graphically as well as a command line interface. [#fn1 (1)] For the purpose of maintaining disks we only need ssaccli and perhaps ssaduccli.

- install the hp supported utilities.

```
root@kb2018:~# echo "deb http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current non-free" >> /etc/
apt/sources.list.d/hp.list root@kb2018:~# root@kb2018:/etc/apt# wget http://downloads.linear.hpe.com/SDR/repo/mcp/GPG-
KEY-mcp -2018-11-12 09:00:29-- http://downloads.linear.hpe.com/SDR/repo/mcp/GPG-KEY-mcp Resolving
downloads.linear.hpe.com (downloads.linear.hpe.com)... 15.249.152.85 Connecting to downloads.linear.hpe.com
(downloads.linear.hpe.com)[15.249.152.85]:80... connected. HTTP request sent, awaiting response... 200 OK Length: 994
Saving to: 'GPG-KEY-mcp'
```

```
GPG-KEY-mcp 100%
```

```
[=====
994 --KB/s in 0s
```

```
2018-11-12 09:00:30 (90.5 MB/s) - 'GPG-KEY-mcp' saved [994/994]
```

```
root@kb2018:/etc/apt# apt-key add GPG-KEY-mcp OK root@kb2018:/etc/apt# apt-get update Ign:1 http://
downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current InRelease Get:2 http://security.ubuntu.com/ubuntu
bionic-security InRelease [83.2 kB]
```

```
Hit:3 http://us.archive.ubuntu.com/ubuntu bionic InRelease
```

```
Get:4 http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release [6,051 B]
```

```
Get:5 http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release.gpg [490 B]
```

```
Get:6 http://us.archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
```

```
Hit:7 http://archive.ubuntu.com/ubuntu bionic InRelease
```

```
Ign:5 http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release.gpg
```

```
Get:8 http://archive.ubuntu.com/ubuntu bionic-security InRelease [83.2 kB]
```

```
Get:9 http://archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
```

```
Reading package lists... Done
```

```
W: GPG error: http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release: The following signatures
couldn't be verified because the public key is not available: NO_PUBKEY C208ADDE26C2B797 E: The repository 'http://
downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release' is not signed. N: Updating from such a
repository can't be done securely, and is therefore disabled by default. N: See apt-secure(8) manpage for repository creation
and user configuration details. root@kb2018:/etc/apt# key=C208ADDE26C2B797 root@kb2018:/etc/apt# gpg --keyserver
keyserver.ubuntu.com --recv-keys $key gpg: key C208ADDE26C2B797: public key "Hewlett Packard Enterprise Company
RSA-2048-25 signhp@hpe.com" imported gpg: Total number processed: 1 gpg: imported: 1 root@kb2018:/etc/apt# gpg --
armor --export $key [apt-key add - OK root@kb2018:/etc/apt# apt-get update Get:1 http://security.ubuntu.com/ubuntu
bionic-security InRelease [83.2 kB] Hit:2 http://us.archive.ubuntu.com/ubuntu bionic InRelease
Get:3 http://us.archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Hit:4 http://archive.ubuntu.com/ubuntu bionic InRelease
Ign:5 http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current InRelease
Get:6 http://archive.ubuntu.com/ubuntu bionic-security InRelease [83.2 kB]
Get:7 http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release [6,051 B]
Get:8 http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release.gpg [490 B] Get:9 http://
archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Get:10 http://downloads.linear.hpe.com/SDR/downloads/MCP/ubuntu bionic/current/non-free amd64 Packages [1,971 B]
Fetched 352 kB in 1s (288 kB/s)
Reading package lists... Done root@kb2018:/etc/apt# apt-get install ssaccli ssaduccli ...
```

Once the issues with his signature were resolved (above) I was able to instal the ssaccli. [#fn2 (2)]

SEEING THE DRIVES

Use the ssaccli to show the unassigned drives after inserting fresh disks.

```
root@kb2018:/etc/apt# ssaccli
Smart Storage Administrator CLI 3.30.13.0
Detecting Controllers...Done.
Type "help" for a list of supported commands.
Type "exit" to close the console.
```

```
=> set target controller slot=0

"controller slot=0"

=> pd all show

Smart Array P410i in Slot 0 (Embedded)

  Array A

    physicaldrive 2C:1:1 (port 2C:box 1:bay 1, SAS HDD, 146 GB, OK)
    physicaldrive 2C:1:2 (port 2C:box 1:bay 2, SAS HDD, 146 GB, OK)

  Array B

    physicaldrive 2C:1:3 (port 2C:box 1:bay 3, SATA SSD, 240 GB, OK)

  Array C

    physicaldrive 2C:1:4 (port 2C:box 1:bay 4, SATA SSD, 240 GB, OK)

  Array D

    physicaldrive 3C:1:5 (port 3C:box 1:bay 5, SAS HDD, 300 GB, OK)

  Array E

    physicaldrive 3C:1:6 (port 3C:box 1:bay 6, SAS HDD, 300 GB, OK)

  Unassigned

    physicaldrive 3C:1:7 (port 3C:box 1:bay 7, SAS HDD, 146 GB, OK)
    physicaldrive 3C:1:8 (port 3C:box 1:bay 8, SAS HDD, 146 GB, OK)
```

LETTING THE OS SEE THE DRIVES

Once we know what drives are available we can create logical drives which will be presented to the os (*assuming the same set target command above*)

```
=> set target controller slot=0
...
=> create type=ld drives=3C:1:7 size=max raid=0
=> create type=ld drives=3C:1:8 size=max raid=0
quit
```

PREPARING THE DRIVES FOR REMOVAL

Before removing a drive you should delete the logical disk that it is associated with.

```
=> set target controller slot=0
...
=> Array G delete
```

INCREASING WRITE PERFORMANCE

once we get a ups we should be able to use the controllers write cache safely.

```
=> controller slot=0 modify drivewritecache=enable

Warning: Without the proper safety precautions, use of write cache on physical
drives could cause data loss in the event of power failure. To ensure
data is properly protected, use redundant power supplies and
Uninterruptible Power Supplies. Also, if you have multiple storage
enclosures, all data should be mirrored across them. Use of this
feature is not recommended unless these precautions are followed.
Continue? (y/n) n

=>
```

footnotes

[=#fn1 1]) This was discovered after digging around for the perccli raid utilities provided by dell (officially supported only on commercial RPM based systems but installable using alien)

[=#fn2 2]) The biggest pain in the ass other than the weirdness with the public signature was that HP fucking rebranded the hpssacli to ssaccli. Most of the good web info and hp docs still reference the old utility name (nothing else changed).

references

- <http://h10032.www1.hp.com/ctg/Manual/c02289065.pdf> (2010)
- <https://amk1.wordpress.com/2013/11/22/zfs-with-hp-smart-array-p410i/>
- <https://content.etilize.com/User-Manual/1033728289.pdf>
- <http://www.sysadminshare.com/2012/05/hpacucli-commands-reference.html>
- <https://wiki.debian.org/LinuxRaidForAdmins>
- <https://www.golinuxhub.com/2017/05/hot-swapping-broken-hdd-with-software.html>
- <https://kallesplayground.wordpress.com/useful-stuff/hp-smart-array-cli-commands-under-esxi/>
- <http://downloads.linux.hpe.com/SDR/project/mcp/>
- https://wiki.debian.org/HP/ProLiant#HP_Repository
- <https://binaryimpulse.com/2013/09/hp-array-configuration-utility-command-cheat-sheet/>
- <https://bibszone.wordpress.com/2016/02/11/hp-smart-array-cli-commands/>
- https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/doc/general/mgmt/ssa_cli/files/v240_130/hpsacli-2.40-13.0_help.txt
- <https://unixlab.weebly.com/raid-array.html>
- <https://hardforum.com/threads/hp-dl380p-gen8-p420i-controller-hbamode.1852528/>

addendum (output from ssacli show detailed config)

```
=>ctrl all show config detail

Smart Array P410i in Slot 0 (Embedded)
  Bus Interface: PCI
  Slot: 0
  Serial Number: 5001438013631A40
  Cache Serial Number: PBCDH0CRH0V0L0
  Controller Status: OK
  Hardware Revision: C
  Firmware Version: 6.64-0
  Rebuild Priority: Medium
  Expand Priority: Medium
  Surface Scan Delay: 15 secs
  Surface Scan Mode: Idle
  Parallel Surface Scan Supported: No
  Queue Depth: Automatic
  Monitor and Performance Delay: 60 min
  Elevator Sort: Enabled
  Degraded Performance Optimization: Disabled
  Wait for Cache Room: Disabled
  Surface Analysis Inconsistency Notification: Disabled
  Post Prompt Timeout: 0 secs
  Cache Board Present: True
  Cache Status: OK
  Cache Ratio: 25% Read / 75% Write
  Drive Write Cache: Disabled
  Total Cache Size: 0.5
  Total Cache Memory Available: 0.4
  No-Battery Write Cache: Disabled
  Cache Backup Power Source: Capacitors
  Battery/Capacitor Count: 1
  Battery/Capacitor Status: OK
  SATA NCQ Supported: True
  Number of Ports: 2 Internal only
  Encryption: Not Set
  Driver Name: hpsa
  Driver Version: 3.4.20
  Driver Supports SSD Smart Path: True
  PCI Address (Domain:Bus:Device.Function): 0000:05:00.0
  Port Max Phy Rate Limiting Supported: False
  Host Serial Number: USE135N52V
  Sanitize Erase Supported: False
  Primary Boot Volume: None
  Secondary Boot Volume: None

HP SAS Expander Card at Port 2C, Box 1, OK

  Power Supply Status: Not Redundant
  Vendor ID: HP
  Serial Number: RF15BP2689
  Firmware Version: 2.10
  Drive Bays: 24
```

```

Port: 2C
Box: 1
Location: Internal

Expander 250
Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66
Box: 1
Vendor ID: HP

HP SAS Expander Card SEP 248
Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP
Model: HP SAS EXP Card

Physical Drives
physicaldrive 2C:1:4 (port 2C:box 1:bay 4, SATA SSD, 240 GB, OK)
physicaldrive 2C:1:3 (port 2C:box 1:bay 3, SATA SSD, 240 GB, OK)
physicaldrive 2C:1:2 (port 2C:box 1:bay 2, SAS HDD, 146 GB, OK)
physicaldrive 2C:1:1 (port 2C:box 1:bay 1, SAS HDD, 146 GB, OK)
physicaldrive 3C:1:6 (port 3C:box 1:bay 6, SAS HDD, 300 GB, OK)
physicaldrive 3C:1:5 (port 3C:box 1:bay 5, SAS HDD, 300 GB, OK)


HP SAS Expander Card at Port 4C, Box 2, OK

Power Supply Status: Not Redundant
Vendor ID: HP
Serial Number: RF15BP2689
Firmware Version: 2.10
Drive Bays: 24
Port: 4C
Box: 2
Location: Internal

Expander 250
Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66
Box: 1
Vendor ID: HP

HP SAS Expander Card SEP 248
Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP
Model: HP SAS EXP Card

Physical Drives
None attached

Port Name: 1I
Port ID: 0
Port Connection Number: 0
SAS Address: 5001438013631A40
Port Location: Internal

Port Name: 2I
Port ID: 1
Port Connection Number: 1
SAS Address: 5001438013631A44
Port Location: Internal

Array: A
Interface Type: SAS
Unused Space: 6 MB (0.00%)
Used Space: 273.40 GB (100.00%)
Status: OK
Array Type: Data
Smart Path: disable

Logical Drive: 1
Size: 136.70 GB
Fault Tolerance: 1
Heads: 255
Sectors Per Track: 32
Cylinders: 35132
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Unrecoverable Media Errors: None
Caching: Enabled
Unique Identifier: 600508B1001CAA24339C082CBF1B0912

```



```

Disk Name: /dev/sda
Mount Points: / 80.0 GB Partition Number 2
OS Status: LOCKED
Logical Drive Label: A0E0B9A75001438013631A40256F
Mirror Group 1:
    physicaldrive 2C:1:2 (port 2C:box 1:bay 2, SAS HDD, 146 GB, OK)
Mirror Group 2:
    physicaldrive 2C:1:1 (port 2C:box 1:bay 1, SAS HDD, 146 GB, OK)
Drive Type: Data
LD Acceleration Method: Controller Cache

```

```

physicaldrive 2C:1:1
Port: 2C
Box: 1
Bay: 1
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 146 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 15000
Firmware Revision: HPDD
Serial Number: PLWGTWSE
WWID: 5000CCA00B53489D
Model: HP EH0146FARMD
Current Temperature (C): 35
Maximum Temperature (C): 42
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

```

physicaldrive 2C:1:2
Port: 2C
Box: 1
Bay: 2
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 146 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 15000
Firmware Revision: HPDD
Serial Number: PLWP0XNE
WWID: 5000CCA00B5E9B11
Model: HP EH0146FARMD
Current Temperature (C): 34
Maximum Temperature (C): 47
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

```

Array: B
Interface Type: Solid State SATA
Unused Space: 2 MB (0.00%)
Used Space: 223.54 GB (100.00%)
Status: OK
Array Type: Data
Smart Path: disable

```

```

Logical Drive: 2
Size: 223.54 GB
Fault Tolerance: 0
Heads: 255
Sectors Per Track: 32
Cylinders: 57450
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Caching: Enabled
Unique Identifier: 600508B1001CC841DD71B0E330404FF4
Disk Name: /dev/sdb
Mount Points: None
Logical Drive Label: ABABB8965001438013631A40D1E0
Drive Type: Data
LD Acceleration Method: Controller Cache

```

```

physicaldrive 2C:1:3
Port: 2C
Box: 1
Bay: 3
Status: OK
Drive Type: Data Drive
Interface Type: Solid State SATA
Size: 240 GB
Drive exposed to OS: False

```

```

Logical/Physical Block Size: 512/512
Firmware Revision: Q0410A
Serial Number: AB20180827A0101371
WWID: 5001438014526C41
Model: ATA      TEAML5Lite3D240G
SATA NCQ Capable: True
SATA NCQ Enabled: True
SSD Smart Trip Wearout: Not Supported
PHY Count: 1
PHY Transfer Rate: 3.0Gbps
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

Array: C

```

Interface Type: Solid State SATA
Unused Space: 2 MB (0.00%)
Used Space: 223.54 GB (100.00%)
Status: OK
Array Type: Data
Smart Path: disable

```

Logical Drive: 3

```

Size: 223.54 GB
Fault Tolerance: 0
Heads: 255
Sectors Per Track: 32
Cylinders: 57450
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Caching: Enabled
Unique Identifier: 600508B1001CD1056D9358D036DE54EB
Disk Name: /dev/sdc
Mount Points: None
Logical Drive Label: ABAB89005001438013631A4045F6
Drive Type: Data
LD Acceleration Method: Controller Cache

```

physicaldrive 2C:1:4

```

Port: 2C
Box: 1
Bay: 4
Status: OK
Drive Type: Data Drive
Interface Type: Solid State SATA
Size: 240 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Firmware Revision: Q0410A
Serial Number: AB20180827A0100293
WWID: 5001438014526C40
Model: ATA      TEAML5Lite3D240G
SATA NCQ Capable: True
SATA NCQ Enabled: True
SSD Smart Trip Wearout: Not Supported
PHY Count: 1
PHY Transfer Rate: 3.0Gbps
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

Array: D

```

Interface Type: SAS
Unused Space: 0 MB (0.00%)
Used Space: 279.37 GB (100.00%)
Status: OK
Array Type: Data
Smart Path: disable

```

Logical Drive: 4

```

Size: 279.37 GB
Fault Tolerance: 0
Heads: 255
Sectors Per Track: 32
Cylinders: 65535
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Caching: Enabled
Unique Identifier: 600508B1001C868C26439B5D426224F
Disk Name: /dev/sdd
Mount Points: None
Logical Drive Label: ABAB99875001438013631A40A72E
Drive Type: Data
LD Acceleration Method: Controller Cache

```

physicaldrive 3C:1:5

```

Port: 3C
Box: 1
Bay: 5
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 300 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 10000
Firmware Revision: HPD6 (FW update is recommended to minimum version: HPD7)
Serial Number: PQJ0EM4B
WWID: 5000CCA025718881
Model: HP      EG0300FBDBR
Current Temperature (C): 31
Maximum Temperature (C): 44
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

Array: E

```

Interface Type: SAS
Unused Space: 0 MB (0.00%)
Used Space: 279.37 GB (100.00%)
Status: OK
Array Type: Data
Smart Path: disable

```

```

Logical Drive: 5
Size: 279.37 GB
Fault Tolerance: 0
Heads: 255
Sectors Per Track: 32
Cylinders: 65535
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Caching: Enabled
Unique Identifier: 600508B1001C380646CF15536E61E692
Disk Name: /dev/sde
Mount Points: None
Logical Drive Label: ABABE9D05001438013631A4088C8
Drive Type: Data
LD Acceleration Method: Controller Cache

```

physicaldrive 3C:1:6

```

Port: 3C
Box: 1
Bay: 6
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 300 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 10000
Firmware Revision: HPD6 (FW update is recommended to minimum version: HPD7)
Serial Number: PMVJ07DB
WWID: 5000CCA0211D1B55
Model: HP      EG0300FBDBR
Current Temperature (C): 31
Maximum Temperature (C): 57
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

Expander 250

```

Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66
Box: 1
Vendor ID: HP

```

Expander 250

```

Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66
Box: 1
Vendor ID: HP

```

HP SAS Expander Card SEP 248

```

Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP

```

```
Model: HP SAS EXP Card

HP SAS Expander Card SEP 248
Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP
Model: HP SAS EXP Card

SEP (Vendor ID PMCSIERA, Model SRC 8x6G) 249
Device Number: 249
Firmware Version: RevC
WWID: 5001438013631A4F
Vendor ID: PMCSIERA
Model: SRC 8x6G
```

5.13 DeeDee

5.13.1 Basement Server Setup

"Hey Daddy O. I don't wanna go, down to the basement"

DeeDee, [wiki:Joey Joey], and [wiki:Annie Annie] are hp z400s intended to be used at the home lans. They provide the following services to the lan. * Dns filtering via pihole * Http/s caching via squid. * Distributed file sharing / private cloud backup (mechanism tbd) * LXD Container based services * Zero Configuration Networking * ZFS/Mirrored File Sharing. * (planned : sso)

Example Home Network Configuration

[[Image(wiki:Annie:Home Network Diagram.jpg, width=70%)]]

INITIAL PDX NETWORK CONFIGURATION

Our portland location has internet through Century Link. Largely because they have proven them selves historically trustable to not sell all of our personable data to the government.

Initially we set up our home servers on a single /24 network. However our router did some things that made me uncomfortable. * The routers operating system is proprietary and can not be replaced (easily) * The default settings were way insecure. * Even though it was a private network the router refused to let it be larger than a class-c network. * The dns from the router provided some sketch redirection including outside resolution of .lan and .local addresses.

IMPROVED PDX NETWORK CONFIGURATION.

Treating the initial network with the same disdain and suspicion as the greater internet beyond we segmented the network into the original private c block, and a half b class network connected via [wiki:GoldCoastRouter a router running openwrt 19.07]. * Local addresses and name resolution are handled by the router. * External DNS is tied to a filtering server hosted on the Home server (Pihole via an LXD Container) * A caching server (squid) is also hosted on one of the servers. * Services which are meant to interact with the outside world are connected to the upstream router treating its network as a DMZ.

Basement Server Hardware.

Our current platform for the home server is the HP z400 workstation which is capable of file serving as well as LXD based containers. All servers should have at least 12G of memory and an sas raid controller, more info on upgrades etc can be found on my [wiki:NotesForHPZ400Workstation Z400 notes page], I also put an ssd and a second GB Nic into Joey for transferring data between the internet, our internet deployed containers, and our home servers.

To stage DeeDee I used one of the 500GB hp disks for booting and one 600G SAS drive for infrastructure containers and data. These disks and the raid controller were sent to be installed once I have the rudimentary system in place. Once tested I recommend installing a second sas drive to mirror the data.

OS Installation.

Much like the servers at the cool Ubuntu 18.04 was installed using the alternate installer, (*Tasksel: Samba, SSH and Basic ubuntu servers*). In addition, zfsutils were installed: not much else.

LXD Configuration

lxd was initialized using the scsi id of the SAS disk (in hopes that the disk will just show up when installed in new system)

```
oot@DeeDee:~# zpool status
pool: infra
state: ONLINE
scan: none requested
config:

NAME                                STATE  READ WRITE CKSUM
infra                                ONLINE    0    0    0
scsi-3600508b1001cd7e650c500a2e7a5a52d  ONLINE    0    0    0
```

Since we have an existing lxd server we allow connections to the daemon.

```
root@DeeDee:~# lxc config set core.https_address [::]:8443
root@DeeDee:~# lxc config set core.trust_password ~something secure~
```

We have a working pi-hole container. Copy it from Annie. _Also the susdev19 profile contains users and some minor tweaks. _

```
root@annie:~# lxc profile copy susdev19 deede:
root@annie:~# lxc snapshot deniro pihole27jul19
root@annie:~# lxc move deniro/pihole27jul19 deede:pihole
```

squid container

To create the container we set up a profile for the disk and network and create it from a ubuntu-lts image.

```
root@DeeDee:~# lxc image copy ubuntu:18.04 local: --alias=ubuntu-lts
root@DeeDee:~# lxc profile create infra
Profile infra created
root@DeeDee:~# lxc profile edit infra
config: {}
description: LXD profile for infrastructure
devices:
  eth0:
    name: eth0
    nictype: bridged
    parent: br0
    type: nic
  root:
    path: /
    pool: infra
    type: disk
name: infra
root@DeeDee:~# lxc init ubuntu-lts squid -p susdev19 -p infra
Creating squid
root@DeeDee:~# lxc start squid
root@DeeDee:~# lxc exec squid bash
root@squid:~# nano /etc/netplan/50-cloud-init.yaml
root@squid:~# reboot
root@squid:~# root@DeeDee:~#
root@DeeDee:~# lxc list
```

NAME	STATE	IPV4	IPV6	TYPE	SNAPSHOTS
pihole	RUNNING	192.168.0.254 (eth0)		PERSISTENT	0
squid	RUNNING	192.168.0.252 (eth0)		PERSISTENT	0

Then we can update the image and install squid.

```
root@DeeDee:~# lxc exec squid bash
root@squid:~# update.sh
----- begin updating squid -----
...
##### done#####
root@squid:~# apt-get install squid
...
Do you want to continue? [Y/n]
...
root@squid:~# nano /etc/squid/squid.conf
acl SSL_ports port 443
acl Safe_ports port 80      # http
acl Safe_ports port 21      # ftp
acl Safe_ports port 443     # https
acl Safe_ports port 70      # gopher
acl Safe_ports port 210     # wais
acl Safe_ports port 1025-65535 # unregistered ports
acl Safe_ports port 280     # http-mgmt
acl Safe_ports port 488     # gss-http
acl Safe_ports port 591     # filemaker
acl Safe_ports port 777     # multiling http
acl CONNECT method CONNECT
http_access deny !Safe_ports
http_access deny CONNECT !SSL_ports
http_access allow localhost manager
http_access deny manager
http_access allow localhost
acl my_internal_net src 192.168.0.0/24
http_access allow my_internal_net
#http_port 3128 transparent
http_port 3128
coredump_dir /var/spool/squid
refresh_pattern ^ftp:      1440  20%  10080
refresh_pattern ^gopher:  1440   0%  1440
refresh_pattern -i (/cgi-bin/|\?) 0    0%  0
```

```
refresh_pattern (Release|Packages(.gz)*)$ 0 20% 2880
refresh_pattern . 0 20% 4320
```

Segmenting the network

I purchased an Asus RT-N56U a while back because it had plenty of memory making it ideal to run openwrt. Once I was able to get a stock 18.06 image on it I set up the new network. The wan interface was set to get its address from the upstream router.

```
root@mullein:/etc/config# nano network
...
config interface 'lan'
    option type 'bridge'
    option ifname 'eth0.1'
    option proto 'static'
    option ipaddr '192.168.129.1'
    option netmask '255.255.128.0'

config interface 'wan'
    option ifname 'eth0.2'
    option proto 'dhcp'
...
```

ADDING A SECOND BRIDGE

By adding a second 1G network card to each of the basement servers we can redefine our network so that the containers face the DMZ while the local (file server/pihole/etc) interface is on the new home network.

```
root@annie# nano /etc/netplan/50-cloud-init.yaml
network:
  version: 2
  renderer: networkd
  ethernets:
    ens6:
      dhcp4: no
      dhcp6: no
    enpls0:
      dhcp4: no
      dhcp6: no
  bridges:
    br1:
      dhcp4: no
      dhcp6: no
      addresses:
        - 192.168.129.69/17
      gateway4: 192.168.129.1
      nameservers:
        addresses:
          - 192.168.129.1
          - 198.202.31.141
      interfaces:
        - enpls0
    br0:
      dhcp4: no
      dhcp6: no
      interfaces:
        - ens6
root@annie # netplan apply
```

.... *todo: Document moving containers to br1*

5.13.2 References / Notes

YAML FROM LXD INIT

```
config: {}
networks: []
storage_pools:
- config:
    source: /dev/disk/by-id/scsi-3600508b1001cd7e650c500a2e7a5a52d
    description: ""
    name: infra
    driver: zfs
profiles:
- config: {}
  description: ""
  devices:
    eth0:
      name: eth0
      nictype: bridged
      parent: br0
      type: nic
  root:
```

```
path: /  
pool: infra  
type: disk  
name: default  
cluster: null
```


5.14 BS2020 LXC to LXD Notes

When I set up BS2020 a year ago I was new to LXC and LXD (both of which are used at present because I wanted separate disk pools and network for infrastructure and development(/deployment). I believe that once we move from 16.04 to 18.04 we should be able to remove LXC from the equation. Regardless the initial pools were set up using lxd init as described in [the install notes for BS2020](#).

```
root@bs2020:~# lxd init
Name of the storage backend to use (dir or zfs) [default=zfs]:
Create a new ZFS pool (yes/no) [default=yes]? yes
Name of the new ZFS pool [default=lxd]: lxd4infra
Would you like to use an existing block device (yes/no) [default=no]? yes
Path to the existing block device: /dev/sde1
Would you like LXD to be available over the network (yes/no) [default=no]?
Do you want to configure the LXD bridge (yes/no) [default=yes]? no

...

root@bs2020:~# lxd init
... create new zfs pool and use all of /dev/sdd1 do not configure bridge ...
root@bs2020:~# dpkg-reconfigure -p medium lxd
... no yes br1 ... use existing bridge...
root@bs2020:~#
```

When we moved disks from the original server to the new one the OS renumbered the disks so that the boot disk is at /dev/sdc1 (bay 2?) and our archive disk is in the last bay (bay 5). The remaining disks make up two zfs pools. In hindsight I would have preferred to initialize the original disks as entire disks as apposed to the first slice.

```
root@bs2020:~# df -k
...
/dev/sdc1          41921808 2140004  37629256   6% /
...
/dev/sdf1          480589544 1524692 454629192   1% /archive
...
root@bs2020:~#

root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h6m with 0 errors on Sun May 13 00:30:55 2018
config:

   NAME        STATE      READ WRITE CKSUM
   lxd4dev      ONLINE     0     0     0
     sdd1       ONLINE     0     0     0
     sdb        ONLINE     0     0     0
     sde        ONLINE     0     0     0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: scrub repaired 0 in 0h2m with 0 errors on Sun May 13 00:26:10 2018
config:

   NAME        STATE      READ WRITE CKSUM
   lxd4infra    ONLINE     0     0     0
     sda1       ONLINE     0     0     0

errors: No known data errors
```

Disks sdb and sde were added to the dev/deployment pool as follows.

```
zpool add -f lxd4dev /dev/sde
zpool add -f lxd4dev /dev/sdb
```

As the disks were previously used they should have been wiped first.

```
wipefs -a /dev/sdf
```

5.14.1 New Disks New Errors

Since adding the two new disks I keep getting io errors on one of them. They do not seem to be causing any data errors however.

```
ZFS has detected an io error:

eid: 106
```

```
class: io
host: bs2020
time: 2018-05-31 13:45:36-0700
vtype: disk
vpath: /dev/sdb1
vguid: 0x04009EB732FC8852
cksum: 0
read: 0
write: 0
pool: lxd4dev
```

5.14.2 Backing up containers using zfs

Link dump

- <http://www.digithink.com/serverdocs/FunWithLinuxDisks>
- <https://www.thegeekdiary.com/zfs-tutorials-creating-zfs-snapshot-and-clones/>

5.15 DDRescue Notes

Dropped my laptop on the way to practice. Thanks to el-capitan the handy network backups stopped working. Amongst the things lost were.

Bookmarks Passwords (Paypal,Ebay,TurboTax,upers)

5.15.1 Linkdump

- <https://arstechnica.com/civis/viewtopic.php?t=1197911>
- http://www.gnu.org/software/ddrescue/manual/ddrescue_manual.html#Examples
- <http://www.happymac.info/cms/knowledge-base/tech-advice/103-disc-recovery-using-ddrescue.html>
- <https://ubuntuforums.org/showthread.php?t=1914085>

5.16 Docker Installation on FranklinOnce the LXD container for docker was built out I followed the to install docker-ce

```

root@franklin:~# apt-get remove docker docker-engine docker.io
...
Removing docker (1.5-1) ...
...
root@franklin:~# curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo apt-key add -
OK
root@franklin:~# add-apt-repository \
> "deb [arch=amd64] https://download.docker.com/linux/ubuntu \
> $(lsb_release -cs) \
> stable"
root@franklin:~# apt-get update
... Done
root@franklin:~# apt-get install docker-ce
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  aufs-tools cgroupfs-mount libltdl7
Suggested packages:
  mountall
The following NEW packages will be installed:
  aufs-tools cgroupfs-mount docker-ce libltdl7
0 upgraded, 4 newly installed, 0 to remove and 0 not upgraded.
Need to get 21.2 MB of archives.
After this operation, 100 MB of additional disk space will be used.
Do you want to continue? [Y/n]
....
root@franklin:~# exit
-w, --workdir string          Working directory inside the container
feurig@franklin:~$ sudo docker run hello-world
[sudo] password for feurig:
Unable to find image 'hello-world:latest' locally
latest: Pulling from library/hello-world
5b0f327be733: Pull complete
Digest: sha256:07d5f7800dfe37b8c2196c7b1c524c33808ce2e0f74e7aa00e603295ca9a0972
Status: Downloaded newer image for hello-world:latest

Hello from Docker!
.....

```

5.16.1 References

- <https://docs.docker.com/engine/installation/linux/docker-ce/ubuntu/>
- <https://www.digitalocean.com/community/tutorials/how-to-install-and-use-docker-on-ubuntu-16-04>

5.17 DL380 Raid Notes

5.17.1 Problem: Where are my disks???

When we installed the os on our new (to us) proliant DL380, Only a single disk was visible in spite of there having been 6 disks installed. This is because the DL380s disk controller was set up in raid mode and did not expose disks until they were configured as "logical" disks.

This is unlike the Dell PowerEdge we have which detects and presents the drives in a hot swappable fashion while still allowing some disks to participate in raid arrays.

Since we use hardware raid mirroring on the boot disks, Adding, removing or replacing disks requires configuring the raid controller.

5.17.2 Using HP utilities to configure the controller without downing the server

HP provides utilities and officially supports bionic and hosts a repo for it. It includes a server that can be accessed graphically as well as a command line interface. [1] For the purpose of maintaining disks we only need ssacli and perhaps ssaduccli.

Install the hp supported utilities.

```
root@kb2018:~# echo "deb http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current non-free" >> /etc/apt/sources.list.d/hp.list
root@kb2018:~# root@kb2018:/etc/apt# wget http://downloads.linux.hpe.com/SDR/repo/mcp/GPG-KEY-mcp
--2018-11-12 09:00:29-- http://downloads.linux.hpe.com/SDR/repo/mcp/GPG-KEY-mcp
Resolving downloads.linux.hpe.com (downloads.linux.hpe.com)... 15.249.152.85
Connecting to downloads.linux.hpe.com (downloads.linux.hpe.com)|15.249.152.85|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 994
Saving to: 'GPG-KEY-mcp'

GPG-KEY-mcp                               100%[=====>]          994  --.-KB/s

2018-11-12 09:00:30 (90.5 MB/s) - 'GPG-KEY-mcp' saved [994/994]

root@kb2018:/etc/apt# apt-key add GPG-KEY-mcp
OK
root@kb2018:/etc/apt# apt-get update
Ign:1 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current InRelease
Get:2 http://security.ubuntu.com/ubuntu bionic-security InRelease [83.2 kB]
Hit:3 http://us.archive.ubuntu.com/ubuntu bionic InRelease
Get:4 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release [6,051 B]
Get:5 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release.gpg [490 B]
Get:6 http://us.archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Hit:7 http://archive.ubuntu.com/ubuntu bionic InRelease
Ign:5 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release.gpg
Get:8 http://archive.ubuntu.com/ubuntu bionic-security InRelease [83.2 kB]
Get:9 http://archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Reading package lists... Done
W: GPG error: http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release: The following signatures couldn't be verified because the public key is not available: NO_PUBKEY C208ADDE26C2B797
E: The repository 'http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release' is not signed.
N: Updating from such a repository can't be done securely, and is therefore disabled by default.
N: See apt-secure(8) manpage for repository creation and user configuration details.
root@kb2018:/etc/apt# key=C208ADDE26C2B797
root@kb2018:/etc/apt# gpg --keyserver keyserver.ubuntu.com --recv-keys $key
gpg: key C208ADDE26C2B797: public key "Hewlett Packard Enterprise Company RSA-2048-25 <signhp@hpe.com>" imported
gpg: Total number processed: 1
gpg: imported: 1
root@kb2018:/etc/apt# gpg --armor --export $key |apt-key add -
OK
root@kb2018:/etc/apt# apt-get update
Get:1 http://security.ubuntu.com/ubuntu bionic-security InRelease [83.2 kB]
Hit:2 http://us.archive.ubuntu.com/ubuntu bionic InRelease
Get:3 http://us.archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Hit:4 http://archive.ubuntu.com/ubuntu bionic InRelease
Ign:5 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current InRelease
Get:6 http://archive.ubuntu.com/ubuntu bionic-security InRelease [83.2 kB]
Get:7 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release [6,051 B]
Get:8 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current Release.gpg [490 B]
Get:9 http://archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Get:10 http://downloads.linux.hpe.com/SDR/downloads/MCP/ubuntu bionic/current/non-free amd64 Packages [1,971 B]
Fetched 352 kB in 1s (288 kB/s)
Reading package lists... Done
root@kb2018:/etc/apt# apt-get install ssacli ssaduccli
...
```

Once the issues with his signature were resolved (above) I was able to instal the ssacli. [#fn2 (2)]

SEEING THE DRIVES

Use the ssacli to show the unassigned drives after inserting fresh disks.

```
root@kb2018:/etc/apt# ssacli
Smart Storage Administrator CLI 3.30.13.0
Detecting Controllers...Done.
Type "help" for a list of supported commands.
Type "exit" to close the console.

=> set target controller slot=0

"controller slot=0"

=> pd all show

Smart Array P410i in Slot 0 (Embedded)

  Array A

    physicaldrive 2C:1:1 (port 2C:box 1:bay 1, SAS HDD, 146 GB, OK)
    physicaldrive 2C:1:2 (port 2C:box 1:bay 2, SAS HDD, 146 GB, OK)

  Array B

    physicaldrive 2C:1:3 (port 2C:box 1:bay 3, SATA SSD, 240 GB, OK)

  Array C

    physicaldrive 2C:1:4 (port 2C:box 1:bay 4, SATA SSD, 240 GB, OK)

  Array D

    physicaldrive 3C:1:5 (port 3C:box 1:bay 5, SAS HDD, 300 GB, OK)

  Array E

    physicaldrive 3C:1:6 (port 3C:box 1:bay 6, SAS HDD, 300 GB, OK)

  Unassigned

    physicaldrive 3C:1:7 (port 3C:box 1:bay 7, SAS HDD, 146 GB, OK)
    physicaldrive 3C:1:8 (port 3C:box 1:bay 8, SAS HDD, 146 GB, OK)
```

LETTING THE OS SEE THE DRIVES

Once we know what drives are available we can create logical drives which will be presented to the os (*assuming the same set target command above*)

```
=> set target controller slot=0
...
=> create type=ld drives=3C:1:7 size=max raid=0
=> create type=ld drives=3C:1:8 size=max raid=0
quit
```

REMOVING DRIVE

Before removing drives you should make sure that they are unmounted or detached (zfs). After removing a drive you should delete the logical disk that it is associated with.

```
=> set target controller slot=0
...
=> Array G delete
```

INCREASING WRITE PERFORMANCE

once we get a ups we should be able to use the controllers write cache safely.

```
=> controller slot=0 modify drivewritecache=enable

Warning: Without the proper safety precautions, use of write cache on physical
drives could cause data loss in the event of power failure. To ensure
data is properly protected, use redundant power supplies and
Uninterruptible Power Supplies. Also, if you have multiple storage
enclosures, all data should be mirrored across them. Use of this
feature is not recommended unless these precautions are followed.
Continue? (y/n) n

=>
```

See also: [wiki:DL380RaidBios my notes on configuring the disks the hard way]

footnotes

[= #fn1 1]) This was discovered after digging around for the perccli raid utilities provided by dell (officially supported only on commercial RPM based systems but installable using alien)

[= #fn2 2]) The biggest pain in the ass other than the weirdness with the public signature was that HP fucking rebranded the hpssacli to ssaccli. Most of the good web info and hp docs still reference the old utility name (nothing else changed).

references

- <http://h10032.www1.hp.com/ctg/Manual/c02289065.pdf> (2010)
- <https://amk1.wordpress.com/2013/11/22/zfs-with-hp-smart-array-p410i/>
- <https://content.etilize.com/User-Manual/1033728289.pdf>
- <http://www.sysadminshare.com/2012/05/hpacucli-commands-reference.html>
- <https://wiki.debian.org/LinuxRaidForAdmins>
- <https://www.golinuxhub.com/2017/05/hot-swapping-broken-hdd-with-software.html>
- <https://kallesplayground.wordpress.com/useful-stuff/hp-smart-array-cli-commands-under-esxi/>
- <http://downloads.linux.hpe.com/SDR/project/mcp/>
- https://wiki.debian.org/HP/ProLiant#HP_Repository
- <https://binaryimpulse.com/2013/09/hp-array-configuration-utility-command-cheat-sheet/>
- <https://bibszone.wordpress.com/2016/02/11/hp-smart-array-cli-commands/>
- https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/doc/general/mgmt/ssa_cli/files/v240_130/hpssacli-2.40-13.0_help.txt
- <https://unixlab.weebly.com/raid-array.html>
- <https://hardforum.com/threads/hp-dl380p-gen8-p420i-controller-hbamode.1852528/>

addendum (output from ssaccli show detailed config)

```
=>ctrl all show config detail

Smart Array P410i in Slot 0 (Embedded)
  Bus Interface: PCI
  Slot: 0
  Serial Number: 5001438013631A40
  Cache Serial Number: PBCDH0CRH0V0L0
  Controller Status: OK
  Hardware Revision: C
  Firmware Version: 6.64-0
  Rebuild Priority: Medium
  Expand Priority: Medium
  Surface Scan Delay: 15 secs
  Surface Scan Mode: Idle
  Parallel Surface Scan Supported: No
  Queue Depth: Automatic
  Monitor and Performance Delay: 60 min
  Elevator Sort: Enabled
  Degraded Performance Optimization: Disabled
  Wait for Cache Room: Disabled
  Surface Analysis Inconsistency Notification: Disabled
  Post Prompt Timeout: 0 secs
  Cache Board Present: True
  Cache Status: OK
  Cache Ratio: 25% Read / 75% Write
  Drive Write Cache: Disabled
  Total Cache Size: 0.5
  Total Cache Memory Available: 0.4
  No-Battery Write Cache: Disabled
  Cache Backup Power Source: Capacitors
  Battery/Capacitor Count: 1
  Battery/Capacitor Status: OK
  SATA NCQ Supported: True
  Number of Ports: 2 Internal only
  Encryption: Not Set
  Driver Name: hpsa
  Driver Version: 3.4.20
```

```

Driver Supports SSD Smart Path: True
PCI Address (Domain:Bus:Device.Function): 0000:05:00.0
Port Max Phy Rate Limiting Supported: False
Host Serial Number: USE135N52V
Sanitize Erase Supported: False
Primary Boot Volume: None
Secondary Boot Volume: None

```

HP SAS Expander Card at Port 2C, Box 1, OK

```

Power Supply Status: Not Redundant
Vendor ID: HP
Serial Number: RF15BP2689
Firmware Version: 2.10
Drive Bays: 24
Port: 2C
Box: 1
Location: Internal

```

Expander 250

```

Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66
Box: 1
Vendor ID: HP

```

HP SAS Expander Card SEP 248

```

Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP
Model: HP SAS EXP Card

```

Physical Drives

```

physicaldrive 2C:1:4 (port 2C:box 1:bay 4, SATA SSD, 240 GB, OK)
physicaldrive 2C:1:3 (port 2C:box 1:bay 3, SATA SSD, 240 GB, OK)
physicaldrive 2C:1:2 (port 2C:box 1:bay 2, SAS HDD, 146 GB, OK)
physicaldrive 2C:1:1 (port 2C:box 1:bay 1, SAS HDD, 146 GB, OK)
physicaldrive 3C:1:6 (port 3C:box 1:bay 6, SAS HDD, 300 GB, OK)
physicaldrive 3C:1:5 (port 3C:box 1:bay 5, SAS HDD, 300 GB, OK)

```

HP SAS Expander Card at Port 4C, Box 2, OK

```

Power Supply Status: Not Redundant
Vendor ID: HP
Serial Number: RF15BP2689
Firmware Version: 2.10
Drive Bays: 24
Port: 4C
Box: 2
Location: Internal

```

Expander 250

```

Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66
Box: 1
Vendor ID: HP

```

HP SAS Expander Card SEP 248

```

Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP
Model: HP SAS EXP Card

```

Physical Drives

None attached

Port Name: 1I

```

Port ID: 0
Port Connection Number: 0
SAS Address: 5001438013631A40
Port Location: Internal

```

Port Name: 2I

```

Port ID: 1
Port Connection Number: 1
SAS Address: 5001438013631A44
Port Location: Internal

```

Array: A

```

Interface Type: SAS
Unused Space: 6 MB (0.00%)
Used Space: 273.40 GB (100.00%)

```



```

Status: OK
Array Type: Data
Smart Path: disable

Logical Drive: 1
Size: 136.70 GB
Fault Tolerance: 1
Heads: 255
Sectors Per Track: 32
Cylinders: 35132
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Unrecoverable Media Errors: None
Caching: Enabled
Unique Identifier: 600508B1001CAA24339C082CBF1B0912
Disk Name: /dev/sda
Mount Points: / 80.0 GB Partition Number 2
OS Status: LOCKED
Logical Drive Label: A0E0B9A75001438013631A40256F
Mirror Group 1:
    physicaldrive 2C:1:2 (port 2C:box 1:bay 2, SAS HDD, 146 GB, OK)
Mirror Group 2:
    physicaldrive 2C:1:1 (port 2C:box 1:bay 1, SAS HDD, 146 GB, OK)
Drive Type: Data
LD Acceleration Method: Controller Cache

```

```

physicaldrive 2C:1:1
Port: 2C
Box: 1
Bay: 1
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 146 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 15000
Firmware Revision: HPDD
Serial Number: PLWGTWSE
WWID: 5000CCA00B53489D
Model: HP          EH0146FARWD
Current Temperature (C): 35
Maximum Temperature (C): 42
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

```

physicaldrive 2C:1:2
Port: 2C
Box: 1
Bay: 2
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 146 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 15000
Firmware Revision: HPDD
Serial Number: PLWP0XNE
WWID: 5000CCA00B5E9B11
Model: HP          EH0146FARWD
Current Temperature (C): 34
Maximum Temperature (C): 47
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

```

Array: B
Interface Type: Solid State SATA
Unused Space: 2 MB (0.00%)
Used Space: 223.54 GB (100.00%)
Status: OK
Array Type: Data
Smart Path: disable

```

```

Logical Drive: 2
Size: 223.54 GB
Fault Tolerance: 0
Heads: 255
Sectors Per Track: 32
Cylinders: 57450
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Caching: Enabled

```

```

Unique Identifier: 600508B1001CC841DD71B0E330404FF4
Disk Name: /dev/sdb
Mount Points: None
Logical Drive Label: ABABB8965001438013631A40D1E0
Drive Type: Data
LD Acceleration Method: Controller Cache

```

```

physicaldrive 2C:1:3
  Port: 2C
  Box: 1
  Bay: 3
  Status: OK
  Drive Type: Data Drive
  Interface Type: Solid State SATA
  Size: 240 GB
  Drive exposed to OS: False
  Logical/Physical Block Size: 512/512
  Firmware Revision: Q0410A
  Serial Number: AB20180827A0101371
  WWID: 5001438014526C41
  Model: ATA TEAML5Lite3D240G
  SATA NCQ Capable: True
  SATA NCQ Enabled: True
  SSD Smart Trip Wearout: Not Supported
  PHY Count: 1
  PHY Transfer Rate: 3.0Gbps
  Sanitize Erase Supported: False
  Shingled Magnetic Recording Support: None

```

```

Array: C
  Interface Type: Solid State SATA
  Unused Space: 2 MB (0.00%)
  Used Space: 223.54 GB (100.00%)
  Status: OK
  Array Type: Data
  Smart Path: disable

```

```

Logical Drive: 3
  Size: 223.54 GB
  Fault Tolerance: 0
  Heads: 255
  Sectors Per Track: 32
  Cylinders: 57450
  Strip Size: 256 KB
  Full Stripe Size: 256 KB
  Status: OK
  Caching: Enabled
  Unique Identifier: 600508B1001CD1056D9358D036DE54EB
  Disk Name: /dev/sdc
  Mount Points: None
  Logical Drive Label: ABAB89005001438013631A4045F6
  Drive Type: Data
  LD Acceleration Method: Controller Cache

```

```

physicaldrive 2C:1:4
  Port: 2C
  Box: 1
  Bay: 4
  Status: OK
  Drive Type: Data Drive
  Interface Type: Solid State SATA
  Size: 240 GB
  Drive exposed to OS: False
  Logical/Physical Block Size: 512/512
  Firmware Revision: Q0410A
  Serial Number: AB20180827A0100293
  WWID: 5001438014526C40
  Model: ATA TEAML5Lite3D240G
  SATA NCQ Capable: True
  SATA NCQ Enabled: True
  SSD Smart Trip Wearout: Not Supported
  PHY Count: 1
  PHY Transfer Rate: 3.0Gbps
  Sanitize Erase Supported: False
  Shingled Magnetic Recording Support: None

```

```

Array: D
  Interface Type: SAS
  Unused Space: 0 MB (0.00%)
  Used Space: 279.37 GB (100.00%)
  Status: OK
  Array Type: Data
  Smart Path: disable

```

```

Logical Drive: 4
  Size: 279.37 GB

```

```

Fault Tolerance: 0
Heads: 255
Sectors Per Track: 32
Cylinders: 65535
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Caching: Enabled
Unique Identifier: 600508B1001C868C26439B55D426224F
Disk Name: /dev/sdd
Mount Points: None
Logical Drive Label: ABAB99875001438013631A40A72E
Drive Type: Data
LD Acceleration Method: Controller Cache

```

```

physicaldrive 3C:1:5
Port: 3C
Box: 1
Bay: 5
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 300 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 10000
Firmware Revision: HPD6 (FW update is recommended to minimum version: HPD7)
Serial Number: PQJ0EM4B
WWID: 5000CCA025718881
Model: HP      EG0300FDBDR
Current Temperature (C): 31
Maximum Temperature (C): 44
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

```

Array: E
Interface Type: SAS
Unused Space: 0 MB (0.00%)
Used Space: 279.37 GB (100.00%)
Status: OK
Array Type: Data
Smart Path: disable

```

```

Logical Drive: 5
Size: 279.37 GB
Fault Tolerance: 0
Heads: 255
Sectors Per Track: 32
Cylinders: 65535
Strip Size: 256 KB
Full Stripe Size: 256 KB
Status: OK
Caching: Enabled
Unique Identifier: 600508B1001C380646CF15536E61E692
Disk Name: /dev/sde
Mount Points: None
Logical Drive Label: ABABE9D05001438013631A4088C8
Drive Type: Data
LD Acceleration Method: Controller Cache

```

```

physicaldrive 3C:1:6
Port: 3C
Box: 1
Bay: 6
Status: OK
Drive Type: Data Drive
Interface Type: SAS
Size: 300 GB
Drive exposed to OS: False
Logical/Physical Block Size: 512/512
Rotational Speed: 10000
Firmware Revision: HPD6 (FW update is recommended to minimum version: HPD7)
Serial Number: PMVJ07DB
WWID: 5000CCA0211D1855
Model: HP      EG0300FDBDR
Current Temperature (C): 31
Maximum Temperature (C): 57
PHY Count: 2
PHY Transfer Rate: 6.0Gbps, Unknown
Sanitize Erase Supported: False
Shingled Magnetic Recording Support: None

```

```

Expander 250
Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66

```

```
Box: 1
Vendor ID: HP

Expander 250
Device Number: 250
Firmware Version: 2.10
WWID: 5001438014526C66
Box: 1
Vendor ID: HP

HP SAS Expander Card SEP 248
Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP
Model: HP SAS EXP Card

HP SAS Expander Card SEP 248
Device Number: 248
Firmware Version: 2.10
Hardware Revision: Rev C
WWID: 5001438014526C65
Box: 2
Vendor ID: HP
Model: HP SAS EXP Card

SEP (Vendor ID PMCSIERA, Model SRC 8x6G) 249
Device Number: 249
Firmware Version: RevC
WWID: 5001438013631A4F
Vendor ID: PMCSIERA
Model: SRC 8x6G
```

5.18 Esp8266

5.18.1 Linkdump

- <https://learn.sparkfun.com/tutorials/esp8266-thing-development-board-hookup-guide/setting-up-arduino>
- <https://github.com/nodemcu/nodemcu-firmware>
- <https://frightanic.com/iot/comparison-of-esp8266-nodemcu-development-boards/>
- <http://www.esp8266.com/viewtopic.php?f=13&t=2506> (flash size)
- <https://stackoverflow.com/questions/39631011/how-to-determine-flash-size-of-nodemcu>
- <https://hackaday.com/2017/12/28/antenna-alignment-and-hunting-rogue-access-points-with-the-esp8266/#more-286709>
- <http://hackingbeaver.com/?p=957>
- <http://www.esp8266.com/viewtopic.php?f=13&t=3835> (flash size)
- <https://hackaday.io/project/26879-esp8266-controlled-stretch-limousine>
- <https://stackoverflow.com/questions/39631011/how-to-determine-flash-size-of-nodemcu>
- <https://www.allaboutcircuits.com/projects/update-the-firmware-in-your-esp8266-wi-fi-module/> OTA
- <http://www.switchdoc.com/2016/12/iot-esp8266-tutorial-ota-software-updates-arduino-ide/>
- <https://elementztechblog.wordpress.com/2016/06/28/over-the-air-update-for-esp8266-using-arduino-ide/>
- <http://www.whatimade.today/esp8266-easiest-way-to-program-so-far/> Sonoff links
- <https://github.com/arendst/Sonoff-MQTT-OTA>
- <http://www.andremiller.net/content/upgrading-sonoff-wireless-smart-switch-flash-memory-esp8266>
- <https://www.hackster.io/idreams/getting-started-with-sonoff-rf-98a724>
- <http://geek.adachsoft.com/home/article/id/10/n/SONOFF-ESP8266-update-firmware-with-Aduino-IDE/refid/mz>
- <http://randomnerdtutorials.com/reprogram-sonoff-smart-switch-with-web-server/>
- <https://github.com/altelch/SonoffIR>
- <https://tech.scargill.net/itead-slampher-and-sonoff/>
- <https://github.com/arendst/Sonoff-Tasmota-MQTT>
- <https://mosquitto.org/man/mosquitto-conf5.html>
- <https://www.hivemq.com/blog/mqtt-essentials-part-3-client-broker-connection-establishment>
- <https://www.hivemq.com/blog/mqtt-essentials-part-5-mqtt-topics-best-practices>
- <https://hackaday.com/2016/05/09/minimal-mqtt-building-a-broker/>
- <http://www.steves-internet-guide.com/mosquitto-conf-file/>
- <https://jpmens.net/2014/07/03/the-mosquitto-mqtt-broker-gets-websockets-support/>
listener 1883 listener 9001 protocol websockets

5.19 Feurig

5.19.1 To Do List

- make the garage into usable space.
- create infrastructure for home network.
- Deal with legacy crap ([\[\[VideoRanch3d | Videoranch 3d project\]\]](#) / [\[\[VideoRanchToEC2 | Videoranch Website Modernization\]\]](#))
- Flesh out EMS
- create repo mirror
- Link bitbucket code to this tracking system
- Build out generalized hardware.
- Integrate CI/CD with repository.
- Stage Docker / OpenStack Server for deployment.
- Set up Docker
- Set up DevStack
- Implement current server as a "Service"
- Make SuspectDevices Viable again.
- Streamline blog/flickr/facebook as per Jeena's example.
- Clean up [\[\[3DAngstEtsy | 3DAngst Etsy Shop\]\]](#)
- Integrate instagram/flickr

[[wiki:Trac109Blurb](#) old start page]

5.20 FocalNotes

5.20.1 Postgresql (rev from 10 to 12)

This transition is a bit more of a step than usual. * <https://stackoverflow.com/questions/60409585/how-to-upgrade-postgresql-database-from-10-to-12-without-losing-data-for-openpro> * <https://www.issackelly.com/blog/2020/07/06/postgresql-10-to-12-upgrade> * <https://dev.to/rafaelbernard/postgresql-pgupgrade-from-10-to-12-566i> * <https://www.postgresql.org/docs/12/upgrading.html> * <https://dev.to/rafaelbernard/postgresql-pgupgrade-from-10-to-12-566i> * https://rafael.bernard-araujo.com/postgresql-pg_upgrade-from-10-to-12.php

5.21 Gold Coast

Gold Coast (goldcoast.lan) is the house router for portland. The configuration and this doc are at <https://bitbucket.org/houselan/config/src/master/>

5.21.1 LEDE 19.07 on the Ubiquiti ER-lite3

The Ubiquiti EdgeRouter Lite is my new favorite OpenWrt device. It is fast and inexpensive (\$150 new) and the os is on a USB Stick. [[Image(<https://prd-www-cdn.ubnt.com/media/images/product-features/ER-lite-features-UNMS.jpg>)]]

Pros

- 3 independent Gigabit network ports.
- Serial Console
- Cheap and still supported.
- Stock Edge-os would work for most tasks.
- OS on a USB-stick easiest backup and install EVER.
- 512 K of memory.

Cons (some assembly required)

- Because the stock usb stick and (unused) flash is only 4K LEDE considers it a 4K and are threatening to stop producing stock images after 19.07.
- Third party usb sticks take longer to start up than the on board bootloader (U-boot) expects. So a pause and usb reset need to be configured.

5.21.2 How do I get set up?

Building 19.07 for the device

Getting the source. See LEDE documentation for dependencies.

```
feurig@vasily:~$ git clone https://git.openwrt.org/openwrt/openwrt.git
```

Building for the target

```
feurig@vasily:~$ cd openwrt/
feurig@vasily:~/openwrt$ make clean
feurig@vasily:~/openwrt$ git pull
feurig@vasily:~/openwrt$ ./scripts/feeds update -a
feurig@vasily:~/openwrt$ ./scripts/feeds install -a
feurig@vasily:~/openwrt$ make menuconfig
feurig@vasily:~/openwrt$ make -j8 download world
feurig@vasily:~/openwrt$ mv bin/targets/octeon/generic/openwrt-octeon-ubnt_edgerouter-lite-ext4-sysupgrade.tar.gz ~/firmware/
feurig@vasily:~/openwrt$ ./scripts/diffconfig.sh > ../firmware/openwrt-octeon-ubnt_edgerouter-lite-ext4-sysupgrade.diffconfig
```

Deploying the image

Download the image from vasily

```
feurig@colbert:~$ scp feurig@wrt.suspectdevices.com:firmware/openwrt-octeon-ubnt_edgerouter-lite-ext4-sysupgrade.tar.gz .
```

Format the stick with 2 partitions (142M dos and the remaining linux)

```
root@colbert:~$ # fdisk -l
... On our machine, this is our disk ...
Disk /dev/sda: 7.6 GiB, 8166703104 bytes, 15950592 sectors
...
root@colbert:~$ # fdisk /dev/sda
... Partition disk here ...
```



```

root@colbert:~ # fdisk -l
...
Disk /dev/sda: 7.6 GiB, 8166703104 bytes, 15950592 sectors
Disk model: USB 2.0 FD
...
Device      Boot  Start    End Sectors  Size Id Type
/dev/sda1                2048   292863   290816   142M  c W95 FAT32 (LBA)
/dev/sda2    292864  3710975  3418112   1.6G  83 Linux
...
root@colbert:/home/feurig# mkfs.vfat /dev/sda1
root@colbert:/home/feurig# mkfs.ext4 /dev/sda2

```

Copy firmware to usb stick

```

root@colbert:~ # mkdir scratch
root@colbert:~ # cd scratch/
root@colbert:~ # tar -xf ../openwrt-octeon-ubnt_edgerouter-lite-ext4-sysupgrade.tar.gz
root@colbert:~ # mkdir root oroot kernel
root@colbert:~ # mount /dev/sda1 kernel/
root@colbert:~ # mount /dev/sda2 root/
root@colbert:~ # mount sysupgrade-erlite/root oroot -o loop
root@colbert:~ # cp sysupgrade-erlite/kernel kernel/vmlinux.64
root@colbert:~ # md5sum sysupgrade-erlite/kernel | cut -d' ' -f 1 > kernel/vmlinux.64.md5
root@colbert:~ # rsync -aHAX oroot/* root/
root@colbert:~ # umount kernel root oroot
root@colbert:~ # sync

```

Fixing the bootloader for standard USB Sticks.

If the usb stick used takes longer than the stock one to initialize the boot will fail.

```

don$ screen /dev/tty.usbserial 115200
...
U-Boot 1.1.1 (UBNT Build ID: 4670715-gbd7e2d7) (Build time: May 27 2014 - 11:16:22)
.
BIST check passed.
UBNT_E100 r1:2, r2:18, f:4/71, serial #: 802AA84CE978
MPR 13-00318-18
Core clock: 500 MHz, DDR clock: 266 MHz (532 Mhz data rate)
DRAM: 512 MB
Clearing DRAM..... done
Flash: 4 MB
Net: octeth0, octeth1, octeth2
.
USB: (port 0) scanning bus for devices...
      USB device not responding, giving up (status=0)
1 USB Devices found
      scanning bus for storage devices...
No device found. Not initialized?

```

0

Getting the stock boot command

```

Octeon ubnt_e100# printenv
bootdelay=0
baudrate=115200
download_baudrate=115200
nuke_env=protect off $(env_addr) +$(env_size);erase $(env_addr) +$(env_size)
autoload=n
ethact=octeth0
bootcmd=fatload usb 0 $loadaddr vmlinux.64;bootoctlinux $loadaddr coremask=0x3 root=/dev/sda2 rootdelay=15 rw rootsqimg=squashfs.img rootsqwdir=w mtd
...

```

Copy the bootcmd from the existing environment and add a delay and usb reset

```

Octeon ubnt_e100# setenv bootcmd 'sleep 10;usb reset;fatload usb 0 $loadaddr vmlinux.64;bootoctlinux $loadaddr coremask=0x3 root=/dev/sda2 rootdelay=15 rw rootsqimg=squashfs.img rootsqwdir=w mtd'
Octeon ubnt_e100# saveenv
Octeon ubnt_e100# reset

```

5.21.3 Basic LEDE Configuration

- network
- dnsmasq
- firewall
- /etc/ethers

5.21.4 References

Primary

- [OpenWrt Hardware Page](#)
- <https://web.rory.co.nz/2018/02/edgerouter-lite-3-failing-to-boot/>

Link Pile

- <https://community.ui.com/questions/EdgeMax-rescue-kit-now-you-can-reinstall-EdgeOS-from-scratch/58d474b4-604d-48c9-871d-ff44fd9240f3#M12098>
- <https://www.kc8apf.net/2018/01/ubiquiti-edgerouter-lite-usb-surgery/>
- <https://github.com/sowbug/mkeosimg>
- <http://blog.darrenscott.com/2016/09/03/recovering-an-unresponsive-ubiquiti-edgerouter-lite-router/>
- <https://community.ui.com/questions/New-U-Boot-image-for-better-USB-drive-compatibility/c59436cc-dfca-4fab-a923-ba5cdc688a6f?page=2>

5.22 Goodbye Openstack

This does not work. As far as I can tell you can only install devstack on raw hardware and let it install all of its ever moving dependencies. I was able to do this where pike was 6 months ago but not uninstall and reinstall is using the same version.

I don't believe it can trust anything this moving to be sane let alone secure.

FUCK THIS.

Best attempt at lxc in a container

Now we want to see about the devstack container. It should have its own network interface (eno3) and disk /dev/sdf. * use lxd init to create zfs filesystem for devstack At some point we should figure out how to configure both zfs and the appropriate bridge configuration without these two steps.

```
root@bs2020:~# lxd init
...
root@bs2020:~# zfs list
NAME                                USED  AVAIL  REFER  MOUNTPOINT
lxd4devstack                        243M  132G   19K    none
lxd4infra                          200M  132G   19K    none
lxd4infra/naomi                    200M  132G  200M   /var/lib/lxc/naomi/rootfs
root@bs2020:~#
```

<https://docs.openstack.org/devstack/latest/guides/lxc.html>

- setup br1 use dpkg-reconfigure to point the network at br1

```
root@bs2020:~# nano /etc/network/interfaces ... add the following ... auto br1 iface br1 inet static address 0.0.0.0 dns-
nameservers 198.202.31.132 198.202.31.141 8.8.8.8 bridge_ports eno3
```

```
iface eno3 inet manual root@bs2020:~# ifdown br1 && ifup br1 root@bs2020:~# ip a 1: lo: mtu 65536 qdisc noqueue state
UNKNOWN group default qlen 1 ... 12: br1: mtu 1500 qdisc noqueue state UP group default qlen 1000 link/ether
d4:be:d9:ec:ee:d2 brd ff:ff:ff:ff:ff:ff inet6 fe80::d6be:d9ff:feec:eed2/64 scope link valid_lft forever preferred_lft forever
```

- set up lxc config file

```
root@bs2020:~# nano /etc/lxc/devstack.conf
```

5.23 from <https://docs.openstack.org/devstack/latest/guides/lxc.html>

5.24 Permit access to /dev/loop*

```
lxc.cgroup.devices.allow = b 7:* rwm
```

5.25 Setup access to /dev/net/tun and /dev/kvm

```
lxc.mount.entry = /dev/net/tun dev/net/tun none bind,create=file 0 0 lxc.mount.entry = /dev/kvm dev/kvm none
bind,create=file 0 0
```

5.26 Networking

```
lxc.network.type = veth lxc.network.flags = up lxc.network.link = br1 lxc.network.hwaddr = 00:16:3d:xx:xx:xx
```

```
lxc.network.ipv4 = 198.202.31.160/25 lxc.network.ipv4.gateway = 198.202.31.129
```

```
lxc.start.auto = 1 lxc.start.delay = 7 lxc.start.order = 150
```

- create the image

```
root@bs2020:~# lxc-create -n theswedishchef -t ubuntu -f /etc/lxc/devstack.conf -B zfs \
--zfsroot=lx4devstack \ --packages=bsdmainutils,git,nano,ehtables,openvswitch-common
```

- add local admin users, setup network and lockdown ubuntu user.

```
root@bs2020:~# passwd -l ubuntu -R /var/lib/lxc/theswedishchef/rootfs passwd: password expiry information changed.
root@bs2020:~# cd /var/lib/lxc/theswedishchef/rootfs/ root@bs2020:~# cat ~feurig/passed.add>>etc/passwd
root@bs2020:~# cat ~feurig/shadow.add>>etc/shadow root@bs2020:~# tar -xvzf ~feurig/fnj.tgz drwxr-xr-x root/root 0
2017-09-27 17:58 home/ ... home directories for admins mostly for the following file ... -rw-rw-r-- joe/joe 402 2017-09-25
23:51 home/joe/.ssh/authorized_keys root@bs2020:~# cd root@bs2020:~# usermod -R /var/lib/lxc/theswedishchef/rootfs -G
sudo,root joe root@bs2020:~# usermod -R /var/lib/lxc/theswedishchef/rootfs -G sudo,root feurig root@bs2020:~# groupadd -
R /var/lib/lxc/theswedishchef/rootfs -g 1001 feurig root@bs2020:~# groupadd -R /var/lib/lxc/theswedishchef/rootfs -g 1002
feurig root@bs2020:~# groupadd -R /var/lib/lxc/theswedishchef/rootfs -g 1002 joe root@bs2020:~# cat <>/var/lib/lxc/
theswedishchef/rootfs/etc/resolvconf/resolv.conf.d/base dns-nameserver 198.202.31.132 8.8.8.8 nameserver 198.202.31.132
8.8.8.8 eod root@bs2020:~# cat <>/var/lib/lxc/theswedishchef/rootfs/etc/network/interfaces iface eth0 inet static address
198.202.31.160/25 gateway 198.202.31.129 dns-nameservers 198.202.31.132 198.202.31.141 8.8.8.8 dns-search
suspectdevices.com digithink.com eod2
```

- check for ehtables module

```
root@bs2020:~# lsmod |grep ebt ebt_broute 16384 0 ebt_nat 16384 0 ebt_filter 16384 0 ehtables 36864 3
ehtable_broute,ehtable_nat,ehtable_filter x_tables 36864 9
xt_CHECKSUM,ip_tables,xt_tcpudp,ipt_MASQUERADE,xt_conntrack,iptable_filter,ehtables,ipt_REJECT,iptable_mangle
bridge 126976 1 ehtable_broute
```

- run up instance and install devstack

```
root@bs2020:~# lxc-start -n theswedishchef root@bs2020:~# lxc-attach -n theswedishchef root@theswedishchef:~# apt-
get install --reinstall ca-certificates root@theswedishchef:/# useradd -s /bin/bash -d /opt/stack -m stack root@theswedishchef:/
```

```
# echo "stack ALL=(ALL) NOPASSWD: ALL" | sudo tee /etc/sudoers.d/stack stack ALL=(ALL) NOPASSWD: ALL
root@theswedishchef:/# su - stack stack@theswedishchef:~$ git clone https://git.openstack.org/openstack-dev/devstack
Cloning into 'devstack'... .. done. stack@theswedishchef:~$ cd devstack/ stack@theswedishchef:~/devstack$ nano local.conf
[[local|localrc]] ADMIN_PASSWORD=B0rkB0rkB0rk DATABASE_PASSWORD=$ADMIN_PASSWORD
RABBIT_PASSWORD=$ADMIN_PASSWORD SERVICE_PASSWORD=$ADMIN_PASSWORD PUBLIC_INTERFACE=eth0
HOST_IP=127.0.0.1 FLOATING_RANGE=198.202.31.160/28 PUBLIC_NETWORK_GATEWAY=198.202.31.129
Q_FLOATING_ALLOCATION_POOL=start=198.202.31.161,end=192.202.31.173
```

5.27 IPV4_ADDRS_SAFE_TO_USE=172.31.1.0/24

stack@theswedishchef:~/devstack\$./stack.sh ... don't even look at it just walk away

5.27.1 Approaches Attempted

- <https://stgraber.org/2016/10/26/lxd-2-0-lxd-and-openstack-1112/> (lxd fights with yet another fucking automated deployment system (snapd) this is lxd not lxc.... Bottom line snapd and therefore juju wont run in a container on LTS until at least 18.04
- All three "stable" releases. Most of them had issues with different kernel dependencies.

wasted time

- http://blog.decbug.com/openstack_in_lxc/
- <https://bayton.org/docs/linux/lxd/lxd-zfs-and-bridged-networking-on-ubuntu-16-04-lts/>
- <https://www.simpleprecision.com/ubuntu-16-04-lxd-networking-simple-bridge>
- <http://networkstatic.net/installing-openstack-ml2-neutron-plugin-devstack-fedora/>
- <https://blog.scottlowe.org/2012/08/17/installing-kvm-and-open-vswitch-on-ubuntu/>
- <https://docs.docker.com/engine/installation/linux/docker-ce/ubuntu/#prerequisites>
- https://fedoraproject.org/wiki/OpenStack_devstack
- <https://docs.openstack.org/devstack/latest/guides/single-machine.html>
- <https://serenity-networks.com/how-to-install-openstack-ocata-on-a-single-server-using-devstack/>
- <https://linuxcontainers.org/lxd/getting-started-openstack/>
- <https://jujucharms.com/u/openstack-charmers-next/openstack-lxd/>
- <https://stgraber.org/2016/10/26/lxd-2-0-lxd-and-openstack-1112/>
- <https://insights.ubuntu.com/2016/08/15/lunch-learn-with-openstack-containers/>
- <https://help.nextcloud.com/t/install-fails-on-snap-core-2466-mount-unknown-filesystem-squashfs/19251>
- <https://yourcodeway.com/how-to-install-nextcloud-on-ubuntu-16-04>
- <https://askubuntu.com/questions/925391/unknown-filesystem-squashfs-when-trying-to-mount-snap-packages>
- <https://bugs.launchpad.net/snappy/+bug/1628289> dbclinton (dbclin) wrote on 2017-07-26: #36 Just to update my previous comment: poking around Stéphane Graber's blog a bit suggests to me that I really shouldn't expect success with this using less than 16.10.
- <https://askubuntu.com/questions/869792/unit-snap-core-716-mount-has-failed-on-ubuntu-16-04-lts-rootfs-armhf>
-

5.28 Hardening LEDE

```
BusyBox v1.30.1 () built-in shell (ash)
```

OpenWrt 19.07.3, r10663-85e04e9f46

Add packages

In our build `sudo`, `nano`, and `syslog-ng` are included as well as the utilities to work with passwords and groups (`shadow-useradd`, `shadow-groupadd`, `shadow-usermod`) if your build does not you will need to install them.

```
root@OpenWrt:~# opkg update
root@OpenWrt:~# opkg install shadow-useradd shadow-groupadd shadow-usermod
root@OpenWrt:~# opkg install sudo nano svsluo-ng
```

Add Sudo Users

```
root@OpenWrt:~# useradd -c "Joseph Wayne Dumoulin" -m joe -s /bin/ash
root@OpenWrt:~# useradd -c "D Delmar Davis" -m feurig -s /bin/ash
root@OpenWrt:~# groupadd --system sudo
root@OpenWrt:~# usermod -a -G sudo joe
root@OpenWrt:~# usermod -a -G sudo feurig
root@OpenWrt:~# visudo
...
## Uncomment to allow members of group sudo to execute any command
%sudo    ALL=(ALL) ALL
...
root@OpenWrt:~# passwd feurig
root@OpenWrt:~# passwd joe
```

For each user add their authorized ssh keys.

```
sudo -u feurig ash
cd
mkdir .ssh
nano .ssh/authorized_keys
... add keys ...
```

Disable Root Login

Once you are able to log into the router using your ssh keys you should disable root access. The following is recommended but didnt work. *ALWAYS test that you are unable to login as root.*

```
root@openWrt:~# uci set dropbear.@dropbear[0].PasswordAuth="off"
root@openWrt:~# uci set dropbear.@dropbear[0].RootPasswordAuth="off"
root@openWrt:~# uci commit dropbear
root@openWrt:~# reboot
don@annie:~$ ssh root@192.168.128.215
```

```
BusyBox v1.30.1 () built-in shell (ash)
```

```

OpenWrt 19.07.3, r11063-85e04e9f46
root@OpenWrt:~#

```

```
root@OpenWrt:~#
```

Thats worse than ubuntu:ubuntu *Fuck that! Lock the root account and remove dropbears authorized keys.*

```
root@OpenWrt:~# passwd -l root
root@OpenWrt:~# rm /etc/dropbear/authorized_keys
root@OpenWrt:~# ^D
```

```
don@annie:~$ ssh root@192.168.128.215
root@192.168.128.215: Permission denied (publickey).
```

Now the admin users need to log in using their personal ssh keys and escalate privileges using their password.

```
don@annie:~$ ssh feurig@192.168.128.215
[...]
```

PRESERVING USERS HOME DIRECTORIES

In order to maintain the sudo users during upgrades you need to add /home and /etc/sudoers to the /etc/sysupgrade.conf file. The passwd, shadow, group and other files should already be saved by sysupgrade but the home directory is needed for the users .ssh/ authorized keys.

References

- <https://openwrt.org/docs/guide-user/security/secure.access>

5.29 Overview

You may say to yourself "Well, How did I get here?" (*first attempts at setting things up*)

This is my collection of notes on how the servers we own and operate were set up.

5.29.1 medea, BS2020 and other hardware in the ONB building

[Google Doc with map of bresgal/suspect devices IP Addresses](#)

BS2020, (KB2020) Virtualization server(s)

BS2020

BS2020 is an upgrade for bernie with an eye on modern hardware and virtualization. The hardware is a Dell R610 from Server Monkey with 12 processors and 96 Gig of memory.

[BS2020 Install Notes](#)

KB2020

The addition of a second server is planned using similar hardware.

LXC

[wiki:LXDContainersWithProfile Creating Lxd Containers with static ip and admin users]

DOCKER HOST

[wiki:DockeInstallNotes Initial Docker Install on franklin]

VIRTUALIZED MEDEA

For years now Medea (a pile of junk that Joe found) has been providing DNS, Websites and Email to all of our domains. (including this site until recently) These services have been migrated to LXC containers and moved to BS2020.

[Server Migration -- Theory and practice](#)

ADMIN NETWORK (CURRENTLY THE DOT 1)

In order to provide better security for virtual hosts (and to even consider a secure openstack) a separate network for the administrative lan is required. There are several ways that we could achieve this including the vpn connection that Sudti offered to provide, as well as purchasing a dedicated vpn capable firewall.

In a perfect world we should be able to use openVPN to securely connect to the admin network since both openstack and the idrac use several ports to provide GUI and web interfaces and neither can be securely exposed to the internet.

The availability of cheap wifi routers capable of running openwrt or dd-wrt makes it possible to configure and deploy our own vpn tailored to our needs. When we started this process openWRT had still not merged with LEDE (and it still hasn't but all active development is on the LEDE side including a workable openVPN implementation). We have an openWRT (15.04) router in place which has a set of workaround firewall rules to allow us access to bs2020,

We have provisioned an LEDE router (knight) with a working VPN and are waiting for a MOP to swap it out.

OPENWRT

[OpenWRT Setup](#)

REMOTE CONTROL (DELL IDRAC)

5.29.2 Other stuff on the site

- [wiki:7900NWashburne Home Network adventures]
- [wiki:Feurig feurigs todo list.]
- [wiki:Trac109Blurb old start page]

5.29.3 References

- <https://docs.openstack.org/devstack/latest/guides/lxc.html>
- <https://stackoverflow.com/questions/15658932/completely-remove-openstack-from-system-after-installation-from-devstack-script>
- <https://help.ubuntu.com/lts/serverguide/lxc.html>
- <https://stackoverflow.com/questions/24824325/is-there-a-way-to-use-dnsmasq-and-bind-on-the-same-computer>
- <http://www.itzgeek.com/how-tos/linux/ubuntu-how-tos/setup-linux-container-with-lxc-on-ubuntu-16-04-14-04.html>
- <https://www.digitalocean.com/community/tutorials/initial-server-setup-with-ubuntu-16-04>

5.30 Joey Snippet

Most of the information on DeeDee is true for Joey.

Joey's rebuild is detailed at <https://github.com/feurig/edge-server-configuration>

5.31 LXDContainerWithDockerNotes

5.32 FIRST IMPRESSIONS:

Creating LXD Container with Static IP (and Docker Profile) We want to create a docker capable LXD container using an existing bridge with a static ip and zfs. Then we want to install docker and test it. We will make a copy of this container once the admin users have been added so that we wont have to replicate these tasks. Our security model requires ssh keys to log in AND passwords to escalate privileges.

The first thing we learned is that LXC and LXD are pretty different beasts and that while lxc with lxc-templates is a straightforward way to create containers that act a lot like regular old hardware LXD brings on all of its we love the mother fucking cloud baggage. Major differences had to do with user mapping on the containers files created by root on the host were mapped to nobody on the container, making it really difficult to set up home directories etc. (for a workaround to this see <https://stackoverflow.com/questions/33377916/migrating-lxc-to-lxd>) The second was the way that the network is initialized with the assumption that LXD would be providing the bridge and the context.

5.32.1 First Attempt and zfs/bridge setup

- create zfs container and bridge as before

```
root@bs2020:~# lxd init ... create new zfs pool and use all of /dev/sdd1 do not configure bridge ... root@bs2020:~# dpkg-reconfigure -p medium lxd ... no yes br1 ... use existing bridge... root@bs2020:~# lxc launch ubuntu:16.04 franklin -p default -p docker root@bs2020:~# lxc stop franklin root@bs2020:~# passwd -l ubuntu -R /var/lib/lxd/containers/franklin.zfs/rootfs passwd: user 'ubuntu' does not exist root@bs2020:~# cd /var/lib/lxd/containers/franklin.zfs/rootfs/ root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat ~feurig/passwd.add>>etc/passwd root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat ~feurig/shadow.add>>etc/shadow root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# tar -xzvf ~feurig/fnj.tgz home/feurig ... home/joe/.ssh/authorized_keys root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# usermod -R /var/lib/lxd/containers/franklin.zfs/rootfs -G sudo,root joe root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# usermod -R /var/lib/lxd/containers/franklin.zfs/rootfs -G sudo,root feurig root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# groupadd -R /var/lib/lxd/containers/franklin.zfs/rootfs -g 1001 feurig root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# groupadd -R /var/lib/lxd/containers/franklin.zfs/rootfs -g 1002 joe root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat <>/var/lib/lxd/containers/franklin.zfs/rootfs/etc/resolvconf/resolv.conf.d/base

dns-nameserver 198.202.31.132 198.202.31.141 8.8.8.8 nameserver 198.202.31.132 198.202.31.141 8.8.8.8 eod
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# sed -i 's/^iface eth0/#iface eth0/' /var/lib/lxd/containers/franklin.zfs/rootfs/etc/network/interfaces root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat <>/var/lib/lxd/containers/franklin.zfs/rootfs/etc/network/interfaces
iface eth0 inet static address 198.202.31.201/25 gateway 198.202.31.129 dns-nameservers 198.202.31.132 198.202.31.141 8.8.8.8 dns-search suspectdevices.com digithink.com eod2
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# lxc start franklin
```

- try to log in to instance over the network..... FAIL
 - unlike lxc's ubuntu:16.04, lxd's ubuntu:16.04 has all of the cloud cruft . That and all of the modifications to the containers directory was rootsquashed (rendering it useless).
- ```
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# lxc exec franklin bash root@franklin:~# nano /etc/network/interfaces
root@franklin:~# cat /etc/network/interfaces
```

## 5.33 This file describes the network interfaces available on your system

---

## 5.34 and how to activate them. For more information, see interfaces(5).

---

## 5.35 The loopback network interface

---

```
auto lo
iface lo inet loopback
```

## 5.36 Source interfaces

---

### 5.37 Please check /etc/network/interfaces.d before changing this file

---

### 5.38 as interfaces may have been defined in /etc/network/interfaces.d

---

### 5.39 See LP: #1262951

---

```
source /etc/network/interfaces.d/*.cfg iface eth0 inet static address 198.202.31.201/25 gateway 198.202.31.129 dns-
nameservers 198.202.31.132 198.202.31.141 8.8.8.8 dns-search suspectdevices.com digithink.com
```

- *first thought*: remove all of the cloud crap...

```
root@franklin:~# apt-get remove cloud* ... The following packages will be REMOVED: cloud-guest-utils cloud-init cloud-
initramfs-copymods cloud-initramfs-dyn-netconf 0 upgraded, 0 newly installed, 4 to remove and 0 not upgraded. After this
operation, 1682 kB disk space will be freed. Do you want to continue? [Y/n] Y ... root@franklin:~# nano /etc/network/
interfaces ... add auto eth0 root@franklin:~# reboot root@franklin:~# root@bs2020:/var/lib/lxd/containers/franklin.zfs/
rootfs# lxc list +-----+-----+-----+-----+-----+-----+ | NAME | STATE | IPV4 | IPV6 | TYPE | SNAPSHOTS |
+-----+-----+-----+-----+-----+-----+ | franklin | RUNNING | 198.202.31.201 (eth0) | | PERSISTENT | 0 |
+-----+-----+-----+-----+-----+-----+
```

- *second thought*: Fuck that! Make it work!

#### 5.39.1 Second attempt

---

(create LXD profile for suspect devices development).

```
root@bs2020:~# lxc stop franklin
root@bs2020:~# lxc delete franklin
root@bs2020:~# lxc profile create susdev
```

```
root@bs2020:~# lxc profile edit susdev
...
```

- repeat until you have a working system that can be logged into remotely
- create docker container container

```
root@bs2020:~# lxc profile show susdev config: user.network_mode: link-local user.user-data: | #cloud-config timezone:
America/Vancouver users: - name: feurig passwd: "... SUBSTITUTE REAL PASSWORD HASH HERE" gecos: Donald Delmar
Davis ssh-authorized-keys: - ssh-rss ... SUBSTITUTE REAL KEY HERE ... don@viscious groups: sudo,root shell: /bin/bash -
name: joe passwd: "... SUBSTITUTE REALPASSWORD HASH HERE" gecos: Joseph Wayne Dumoulin ssh-authorized-keys: -
ssh-rss ...SUBSTITUTE REAL KEY HERE... jdumoulin@joeslaptop groups: sudo,root shell: /bin/bash manage_resolv_conf: true
resolv_conf: nameservers: ['198.202.31.141', '198.202.31.132', '8.8.8.8'] searchdomains: - suspectdevices.com -
digithink.com domain: suspectdevices.com options: rotate: true timeout: 1 write_files: # Set static IP address could not get
this to work the "right" way - path: /etc/network/interfaces permissions: '0644' owner: root:root content: | auto lo iface lo inet
loopback auto eth0 # change this after first instantiation iface eth0 inet static address 198.202.31.200 broadcast
198.202.31.255 netmask 255.255.255.128 gateway 198.202.31.129 dns-nameservers 198.202.31.141 198.202.31.132
8.8.8.8 runcmd: # sudo needs to be able to resolve itself to authenticate users # and the users are locked by default - sed -i "s/
^127.0.0.1/#127.0.0.1/" /etc/hosts - echo 127.0.0.1 hostname localhost >>/etc/hosts - passwd joe -u - passwd feurig -u
description: Try to create a sane environment for cloud-init based operating systems devices: eth0: name: eth0 nictype:
bridged parent: br1 type: nic name: susdev root@bs2020:~#

root@bs2020:~# lxc list +-----+-----+-----+-----+-----+-----+ | NAME | STATE | IPV4 | IPV6 | TYPE |
SNAPSHOTS | +-----+-----+-----+-----+-----+-----+ | test13 | RUNNING | 198.202.31.200 (eth0) | |
PERSISTENT | 0 | +-----+-----+-----+-----+-----+-----+ root@bs2020:~# lxc init ubuntu:16 franklin -p
susdev -p docker Creating franklin root@bs2020:~# lxc start franklin root@bs2020:~# lxc exec franklin bash
root@franklin:~# tail -2 /etc/shadow feurig:<>:17453:0:99999:7::: joe:<>:17453:0:99999:7::: root@franklin:~# nano /etc/
network/interfaces root@franklin:~# cat /etc/network/interfaces auto lo iface lo inet loopback auto eth0
```

## 5.40 change this after first instantiation

```
iface eth0 inet static address 198.202.31.201 broadcast 198.202.31.255 netmask 255.255.255.128 gateway
198.202.31.129 dns-nameservers 198.202.31.141 198.202.31.132 8.8.8.8 root@franklin:~# cat /etc/hosts
```

## 5.41 127.0.0.1 localhost

## 5.42 The following lines are desirable for IPv6 capable hosts

```
::1 ip6-localhost ip6-loopback fe00::0 ip6-localnet ff00::0 ip6-mcastprefix ff02::1 ip6-allnodes ff02::2 ip6-allrouters ff02::3 ip6-
allhosts 127.0.0.1 franklin localhost root@franklin:~# reboot root@bs2020:~# lxc list +-----+-----+-----+-----+
+-----+-----+ | NAME | STATE | IPV4 | IPV6 | TYPE | SNAPSHOTS | +-----+-----+-----+-----+-----+
franklin | RUNNING | 198.202.31.201 (eth0) | | PERSISTENT | 0 | +-----+-----+-----+-----+-----+ |
test13 | RUNNING | 198.202.31.200 (eth0) | | PERSISTENT | 0 | +-----+-----+-----+-----+-----+
root@bs2020:~#
```

### 5.42.1 References

- <http://www.whiteboardcoder.com/2016/04/cloud-init-nocloud-with-url-for-meta.html>
- <https://stgraber.org/2016/03/11/lxd-2-0-blog-post-series-012/>
- <https://github.com/lxc/lxd/blob/master/doc/cloud-init.md>
- <http://www.mattjarvis.org.uk/post/lxd-openstack-cloudinit-pt1/>
- <https://sdgsystems.com/blog/understanding-and-using-lxc-and-lxd>
- <http://cloudinit.readthedocs.io/en/latest/topics/examples.html>
- <http://cloudinit.readthedocs.io/en/latest/topics/debugging.html>

## 5.43 Imagebuilder notes

### 5.43.1 Building firmware using imagebuilder

In the same page as the binary releases for openwrt/LEDE is the image builder for that architecture for instance at the bottom <https://downloads.lede-project.org/releases/17.01.4/targets/ar71xx/generic/> there is a link [https://downloads.lede-project.org/releases/17.01.4/targets/ar71xx/generic/lede-imagebuilder-17.01.4-ar71xx-generic.Linux-x86\\_64.tar.xz](https://downloads.lede-project.org/releases/17.01.4/targets/ar71xx/generic/lede-imagebuilder-17.01.4-ar71xx-generic.Linux-x86_64.tar.xz). which we untar into /home/openwrt/xx.xx.xx/ on the sandbox.suspectdevices.com container (where xx.xx.xx is the release number. Then we can build the image as follows.

```

root@sandbox:~# cd /home/openwrt/17.01.4/lede-imagebuilder-17.01.4-ar71xx-generic.Linux-x86_64
root@sandbox:/home/openwrt/17.01.4/lede-imagebuilder-17.01.4-ar71xx-generic.Linux-x86_64# make
Available Commands:
 help: This help text
 info: Show a list of available target profiles
 clean: Remove images and temporary build files
 image: Build an image (see below for more information).

Building images:
 By default 'make image' will create an image with the default
 target profile and package set. You can use the following parameters
 to change that:

 make image PROFILE=<profilename> # override the default target profile
 make image PACKAGES=<pkg1> [<pkg2> [<pkg3> ...]] # include extra packages
 make image FILES=<path> # include extra files from <path>
 make image BIN_DIR=<path> # alternative output directory for the images
 make image EXTRA_IMAGE_NAME=<string> # Add this to the output image filename (sanitized)
root@sandbox:/home/openwrt/17.01.4/lede-imagebuilder-17.01.4-ar71xx-generic.Linux-x86_64# make info
Current Target: "ar71xx (Generic)"
Default Packages: base-files libc libgcc busybox dropbear mtd uci pkg netifd fstools uclient-fetch logd kmod-gpio-button-hotplug swconfig kmod-ath9k wpad-mini uboot-
envtools dnsmasq iptables ip6tables ppp ppp-mod-pppoe firewall odhcpd odhcp6c
Available Profiles:

Default:
 Default Profile (all drivers)
 Packages: kmod-usb-core kmod-usb-ohci kmod-usb2 kmod-usb-ledtrig-usbport
ALFAAP120C:
...
wndr3700:
 NETGEAR WNDR3700
 Packages: kmod-usb-core kmod-usb-ohci kmod-usb2 kmod-usb-ledtrig-usbport kmod- leds-wndr3700-usb
wndr3700v2:
 NETGEAR WNDR3700 v2
 Packages: kmod-usb-core kmod-usb-ohci kmod-usb2 kmod-usb-ledtrig-usbport kmod- leds-wndr3700-usb
wndr3800:
...
root@sandbox:/home/openwrt/17.01.4/lede-imagebuilder-17.01.4-ar71xx-generic.Linux-x86_64# make PROFILE="wndr3700v2" PACKAGES="nano" image
make[1]: Entering directory '/home/openwrt/17.01.4/lede-imagebuilder-17.01.4-ar71xx-generic.Linux-x86_64'

```

The resulting firmware will be placed in the bin directory. You can use the factory images to "update" the routers factory firmware to lede. Once you have it installed you can install the next version or future builds using sysupgrade.

[illegible]



```
.uci/
TZ
dhcp.leases
dnsmasq.d/
etc/
hosts/
lede-17.01.4-ar71xx-generic-wndr3700v2-squashfs-sysupgrade.bin
lib/
lock/
log/
overlay/
racoon/
resolv.conf
resolv.conf.auto
run/
state/
sysinfo/
root@vpn:/tmp# sysupgrade -v lede-17.01.4-ar71xx-generic-wndr3700v2-squashfs-sysupgrade.bin
Saving config files...
etc/config/dhcp
etc/config/dropbear
etc/config/firewall
etc/config/luci
etc/config/network
etc/config/rpcd
etc/config/system
etc/config/u-bootenv
etc/config/ucitrack
etc/config/uhttpd
etc/config/wireless
etc/dnsmasq.conf
etc/dropbear/authorized_keys
etc/dropbear/dropbear_dss_host_key
etc/dropbear/dropbear_rsa_host_key
etc/firewall.user
etc/fw_env.config
etc/group
etc/hosts
etc/inittab
etc/iproute2/rt_tables
etc/ipsec.conf
etc/ipsec.secrets
etc/ipsec.user
etc/openssl/ldap.conf
etc/opkg.conf
etc/passwd
etc/ppp/chap-secrets
etc/ppp/filter
etc/ppp/options
etc/ppp/options.xl2tpd
etc/profile
etc/protocols
etc/racoon.conf
etc/racoon/psk.txt
etc/rc.local
etc/services
etc/shadow
etc/shells
etc/ssl/openssl.cnf
etc/strongswan.conf
etc/sysctl.conf
etc/sysupgrade.conf
etc/xl2tpd/xl2tp-secrets
etc/xl2tpd/xl2tpd.conf
killall: watchdog: no process killed
Sending TERM to remaining processes ... odhcpd racoon uhttpd xl2tpd starter charon ntpd odhcp6c dnsmasq ubusd askfirst logd rpcd netifd
Sending KILL to remaining processes ... askfirst
Switching to ramdisk...
Performing system upgrade...
Unlocking firmware ...

Writing from <stdin> to firmware ... [w]
[w]

Appending jffs2 data from /tmp/sysupgrade.tgz to firmware...TRX header not found
Error fixing up TRX header

Upgrade completed
Rebooting system...
```

## References

- <https://openwrt.org/docs/guide-user/additional-software/imagebuilder>
- <http://blog.suspectdevices.com/blahg/electronics/making-due-with-what-you-have/>

## 5.44 OPENVPN on LEDE Notes

Now that we have a recent version of the operating system OpenVPN seems to work as advertised. Following the instructions at <https://lede-project.org/docs/user-guide/openvpn.server>. Much of the heavy lifting is done by easyRSA and MakeOpenVPN.sh.

The client setups fail if you use an empty passphrase which is good. OTOH In my initial attempts I could not get the server certificates to work with one. When in doubt read the documentation sections on the old openWRT site. It provides a little more depth but there still are some missing pieces that require more exploration ([https://wiki.openwrt.org/doc/howto/vpn.openvpn#tab\\_\\_using\\_openssl\\_commands\\_most\\_secure](https://wiki.openwrt.org/doc/howto/vpn.openvpn#tab__using_openssl_commands_most_secure)).

For the client I used tunnelblick which works well and takes the .ovpn configuration files created by this process.

### Sample Install

Follow the bouncing prompt using lede user guide.

```
root@mullein:~# opkg update && opkg install openvpn-openssl openvpn-easy-rsa luci-app-openvpn
Downloading
....note additional dependencies....
Configuring kmod-tun.
Configuring zlib.
Configuring libopenssl.
Configuring openssl-util.
Configuring liblzo.
Configuring openvpn-openssl.
Configuring openvpn-easy-rsa.
Configuring luci-app-openvpn.
root@mullein:~# cd /etc/easy-rsa
root@mullein:/etc/easy-rsa# source vars
NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys
root@mullein:/etc/easy-rsa# clean-all
NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys
root@mullein:/etc/easy-rsa# build-ca
NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'ca.key'

You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.

Country Name (2 letter code) [US]:
State or Province Name (full name) [CA]:OR
Locality Name (eg, city) [SanFrancisco]:Portland
Organization Name (eg, company) [Fort-Funston]:SuspectDevices
Organizational Unit Name (eg, section) [MyOrganizationalUnit]:3dAngst
Common Name (eg, your name or your server's hostname) [Fort-Funston CA]:mullein
Name [EasyRSA]:mullein
Email Address [me@myhost.mydomain]:don@suspectdevices.com
```

Plan on the next step taking so long you will probably have to reconnect and pick up where you were...

```
root@mullein:/etc/easy-rsa# build-dh
NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys
Generating DH parameters, 2048 bit long safe prime, generator 2
This is going to take a long time
....
.... They are not kidding
.....+.....++*++*
```

Continue to follow the bouncing prompt

```
root@mullein:/etc/easy-rsa# build-key-server mullein
.... answer the questions
A challenge password []:
An optional company name []:
Using configuration from /etc/easy-rsa/openssl-1.0.0.cnf
Check that the request matches the signature
Signature ok
The Subject's Distinguished Name is as follows
countryName :PRINTABLE:'US'
stateOrProvinceName :PRINTABLE:'OR'
localityName :PRINTABLE:'Portland'
```

```

organizationName :PRINTABLE:'SuspectDevices'
organizationalUnitName:PRINTABLE:'3dAngst'
commonName :PRINTABLE:'mullein'
name :PRINTABLE:'mullein'
emailAddress :IA5STRING:'don@suspectdevices.com'
Certificate is to be certified until Oct 23 23:46:35 2027 GMT (3650 days)
Sign the certificate? [y/n]:y
1 out of 1 certificate requests certified, commit? [y/n]y
Write out database with 1 new entries
Data Base Updated
root@mullein:/etc/easy-rsa# openvpn --genkey --secret /etc/easy-rsa/keys/ta.key

```

## Set up the network and firewall rules.

```

root@mullein:/etc/easy-rsa# openvpn --genkey --secret /etc/easy-rsa/keys/ta.key
root@mullein:/etc/easy-rsa# uci set network.vpn0="interface"
root@mullein:/etc/easy-rsa# uci set network.vpn0.ifname="tun0"
root@mullein:/etc/easy-rsa# uci set network.vpn0.proto="none"
root@mullein:/etc/easy-rsa# uci set network.vpn0.auto="1"
root@mullein:/etc/easy-rsa# uci commit network
root@mullein:/etc/easy-rsa# uci add firewall rule
cfg1892bd
root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].name="Allow-OpenVPN-Inbound"
root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].target="ACCEPT"
root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].src="wan"
root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].proto="udp"
root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].dest_port="1194"
root@mullein:/etc/easy-rsa# uci add firewall zone
cfg19dc81
root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].name="vpn"
root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].input="ACCEPT"
root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].forward="ACCEPT"
root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].output="ACCEPT"
root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].masq="1"
root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].network="vpn0"
root@mullein:/etc/easy-rsa# uci add firewall forwarding
cfg1aad58
root@mullein:/etc/easy-rsa# uci set firewall.@forwarding[-1].src="vpn"
root@mullein:/etc/easy-rsa# uci set firewall.@forwarding[-1].dest="wan"
root@mullein:/etc/easy-rsa# uci add firewall forwarding
cfg1bad58
root@mullein:/etc/easy-rsa# uci set firewall.@forwarding[-1].src="vpn"
root@mullein:/etc/easy-rsa# uci set firewall.@forwarding[-1].dest="lan"
root@mullein:/etc/easy-rsa# uci commit firewall
root@mullein:/etc/easy-rsa# /etc/init.d/network reload
....
root@mullein:/etc/easy-rsa# /etc/init.d/firewall reload
....

```

## Check ip forwarding

```

root@mullein:/etc/easy-rsa# cat /proc/sys/net/ipv4/ip_forward
1

```

## Edit /etc/config/openvpn, enable and restart daemon.

```

root@mullein:/etc/easy-rsa# nano /etc/config/openvpn
... add the following (change name, cert, and key to match your server) ...
#####
https://lede-project.org/docs/user-guide/openvpn.server
#####
config openvpn 'mullein'
 option enabled '1'
 option dev 'tun'
 option port '1194'
 option proto 'udp'
 option status '/var/log/openvpn_status.log'
 option log '/tmp/openvpn.log'
 option verb '3'
 option mute '5'
 option keepalive '10 120'
 option persist_key '1'
 option persist_tun '1'
 option user 'nobody'
 option group 'nogroup'
 option ca '/etc/easy-rsa/keys/ca.crt'
 option cert '/etc/easy-rsa/keys/mullein.crt'
 option key '/etc/easy-rsa/keys/mullein.key'
 option dh '/etc/easy-rsa/keys/dh2048.pem'
 option mode 'server'
 option tls_server '1'
 option tls_auth '/etc/easy-rsa/keys/ta.key 0'
 option server '10.9.0.0 255.255.255.0'
 option topology 'subnet'
 option route_gateway 'dhcp'
 option client_to_client '1'
 list push 'persist-key'
 list push 'persist-tun'
 list push 'redirect-gateway def1'

```

```
allow your clients to access to your network
list push 'route 192.168.2.0 255.255.255.0'
push DNS to your clients
list push 'dhcp-option DNS 192.168.2.1'
option comp_lzo 'no'

root@mullein:/etc/easy-rsa# /etc/init.d/openvpn start
root@mullein:/etc/easy-rsa# /etc/init.d/openvpn enable
root@mullein:/etc/easy-rsa# cat /tmp/openvpn.log
...
Thu Oct 26 00:22:46 2017 OpenVPN 2.4.3 mipsel-openwrt-linux-gnu [SSL (OpenSSL)] [LZO] [LZ4] [EPOLL] [MH/PKTINFO] [AEAD]
....
Thu Oct 26 00:22:46 2017 MULTI: multi_init called, r=256 v=256
Thu Oct 26 00:22:46 2017 IFCONFIG POOL: base=10.9.0.2 size=252, ipv6=0
Thu Oct 26 00:22:46 2017 Initialization Sequence Completed
...
```

### Create client cert.

```
root@mullein:~# cd /etc/easy-rsa/
root@mullein:/etc/easy-rsa# source vars
NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys
root@mullein:/etc/easy-rsa# build-key-pkcs12 donathome
...
writing new private key to 'donathome.key'
....
Country Name (2 letter code) [US]:
State or Province Name (full name) [CA]:OR
Locality Name (eg, city) [SanFrancisco]:Portland
Organization Name (eg, company) [Fort-Funston]:SuspectDevices
Organizational Unit Name (eg, section) [MyOrganizationalUnit]:3dAngst
Common Name (eg, your name or your server's hostname) [donathome]:viscious
Name [EasyRSA]:DonAtHome
Email Address [me@myhost.mydomain]:don@suspectdevices.com

Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:XXXXXXXXXX
An optional company name []:Its Late
...
Certificate is to be certified until Oct 24 02:49:46 2027 GMT (3650 days)
Sign the certificate? [y/n]:y

1 out of 1 certificate requests certified, commit? [y/n]y
Write out database with 1 new entries
Data Base Updated
Enter Export Password:
Verifying - Enter Export Password:
root@mullein:/etc/easy-rsa# openssl rsa -in /etc/easy-rsa/keys/donathome.key -des3 -out /etc/easy-rsa/keys/donathome.3des.key
writing RSA key
Enter PEM pass phrase:
Verifying - Enter PEM pass phrase:
root@mullein:/etc/easy-rsa#
```

### MakeOpenVPN.sh script (install missing dependencies)

```
root@mullein:/etc/easy-rsa# cd keys
root@mullein:/etc/easy-rsa/keys# wget https://gist.githubusercontent.com/ivanmarban/57561e2bacf3b3a709426d353d2b6584/raw/30bf3c86fbc95a0a5d53d0aac348bcebd9aa2eb/MakeOpenVPN.sh -O /etc/easy-rsa/keys/MakeOpenVPN.sh
wget: SSL support not available, please install one of the libstream-ssl-* libraries as well as the ca-bundle and ca-certificates packages.
root@mullein:/etc/easy-rsa/keys# opkg update && opkg install libstream-openssl ca-certificates
...
root@mullein:/etc/easy-rsa/keys# wget https://gist.githubusercontent.com/ivanmarban/57561e2bacf3b3a709426d353d2b6584/raw/30bf3c86fbc95a0a5d53d0aac348bcebd9aa2eb/MakeOpenVPN.sh -O /etc/easy-rsa/keys/MakeOpenVPN.sh
Downloading 'https://gist.githubusercontent.com/ivanmarban/57561e2bacf3b3a709426d353d2b6584/raw/30bf3c86fbc95a0a5d53d0aac348bcebd9aa2eb/MakeOpenVPN.sh'
Connecting to 151.101.52.133:443
Writing to '/etc/easy-rsa/keys/MakeOpenVPN.sh'
/etc/easy-rsa/keys/M 100% |*****| 1839 0:00:00 ETA
Download completed (1839 bytes)
root@mullein:/etc/easy-rsa/keys# chmod oug+x MakeOpenVPN.sh
```

### Configure and run script.

```
root@mullein:/etc/easy-rsa/keys# nano Default.txt
... Add the following, Adjust host name accordingly
client
dev tun
proto udp
remote mullein.suspectdevices.com 1194
resolv-retry infinite
nobind
mute-replay-warnings
ns-cert-type server
key-direction 1
verb 1
mute 20
comp-lzo no
```

```

root@mullein:/etc/easy-rsa/keys# ./MakeOpenVPN.sh
Please enter an existing Client Name:
donathome
Client's cert found: donathome
Client's Private Key found: donathome.3des.key
CA public Key found: ca.crt
tls-auth Private Key found: ta.key
Done! donathome.ovpn Successfully Created.
root@mullein:/etc/easy-rsa/keys# ls
01.pem ca.crt donathome.key index.txt.old mullein.key myvpn.key
02.pem ca.key donathome.ovpn knight.crt mullien.crt serial
03.pem dh2048.pem donathome.p12 knight.csr mullien.csr serial.old
04.pem donathome.3des.key index.txt knight.key mullien.key ta.key
Default.txt donathome.crt index.txt.attr mullein.crt myvpn.crt
MakeOpenVPN.sh donathome.csr index.txt.attr.old mullein.csr myvpn.csr
root@mullein:/etc/easy-rsa/keys# ./MakeOpenVPN.sh
Please enter an existing Client Name:
donathome
Client's cert found: donathome
Client's Private Key found: donathome.3des.key
CA public Key found: ca.crt
tls-auth Private Key found: ta.key
Done! donathome.ovpn Successfully Created.

```

## References (Link Dump)

- <https://help.my-private-network.co.uk/support/solutions/articles/24000005597-openwrt-lede-openvpn-setup>
- [https://lede-project.org/docs/user-guide/openvpn.server#setup\\_clients](https://lede-project.org/docs/user-guide/openvpn.server#setup_clients)
- <https://steemit.com/openwrt/@rbrthnk/vpn-pptp-router-with-openwrt-lede-tutorial-super-easy>
- [https://lede-project.org/docs/user-guide/tunneling\\_interface\\_protocols](https://lede-project.org/docs/user-guide/tunneling_interface_protocols)
- [https://www.softether.org/4-docs/2-howto/9.L2TPIPsec\\_Setup\\_Guide\\_for\\_SoftEther\\_VPN\\_Server](https://www.softether.org/4-docs/2-howto/9.L2TPIPsec_Setup_Guide_for_SoftEther_VPN_Server)
- [https://wiki.gentoo.org/wiki/IPsec\\_L2TP\\_VPN\\_server](https://wiki.gentoo.org/wiki/IPsec_L2TP_VPN_server)
- [http://connect.rbhs.rutgers.edu/vpn/Mac\\_OSX\\_Native\\_VPN\\_Client\\_Overview.pdf](http://connect.rbhs.rutgers.edu/vpn/Mac_OSX_Native_VPN_Client_Overview.pdf)
- <http://cookbook.fortinet.com/ipsec-vpn-native-mac-os-client-54/>
- <https://www.howtogeek.com/216209/how-to-connect-your-mac-to-any-vpn-and-automatically-reconnect/>
- <https://tunnelblick.net/cInstall.html>
- <https://forum.lede-project.org/t/configuring-lede-router-with-a-pppoe-modem-router/5348/2>
- <https://wiki.openwrt.org/doc/howto/openconnect-setup>
- [https://wiki.gavowen.ninja/doku.php?id=lede:openconnect#tab\\_\\_pki\\_templates](https://wiki.gavowen.ninja/doku.php?id=lede:openconnect#tab__pki_templates)
- <https://lede-project.org/docs/user-guide/openvpn.server>
- [https://wiki.openwrt.org/doc/howto/vpn.openvpn#tab\\_\\_traditional\\_tun\\_client](https://wiki.openwrt.org/doc/howto/vpn.openvpn#tab__traditional_tun_client)

## 5.45 OpenWRT Notes

At a very minimum open the ssh port so that the router can be managed from the outside. Then disable logins (ssh keys only) in /etc/dropbear.

```
root@OpenWrt:/etc/config# opkg update
root@OpenWrt:/etc/config# opkg install nano

root@OpenWrt:/etc/config# nano /etc/config/firewall
```

add the following

```
config redirect
 option target 'DNAT'
 option src 'wan'
 option dest 'lan'
 option proto 'tcp'
 option dest_ip '192.168.1.1'
 option dest_port '22'
 option name 'sshplease'
 option src_dport '2222'
```

### 5.45.1 allowing access to dell IDRAC 6 and server forward

### 5.45.2 firewall setup on vpn

In order to get at the idrac and access BS2020 via ssh the following rules were added to /etc/config/firewall

```
config redirect
 option target 'DNAT'
 option src 'wan'
 option dest 'lan'
 option proto 'tcp'
 option dest_ip '192.168.1.158'
 option dest_port '22'
 option name 'sshtobernie'
 option src_dport '22'

idrac 6 redirections
config redirect
 option target 'DNAT'
 option src 'wan'
 option dest 'lan'
 option proto 'tcp'
 option dest_ip '192.168.1.121'
 option dest_port '443'
 option name 'idracplease1'
 option src_dport '443'

config redirect
 option target 'DNAT'
 option src 'wan'
 option dest 'lan'
 option proto 'tcp'
 option dest_ip '192.168.1.121'
 option dest_port '4433'
 option name 'idracplease2'
 option src_dport '4433'

config redirect
 option target 'DNAT'
 option src 'wan'
 option dest 'lan'
 option proto 'tcp'
 option dest_ip '192.168.1.121'
 option dest_port '443'
 option name 'idracplease3'
 option src_dport '443'

config redirect
 option target 'DNAT'
 option src 'wan'
 option dest 'lan'
 option proto 'tcp'
 option dest_ip '192.168.1.121'
 option dest_port '623'
 option name 'idracplease4'
 option src_dport '623'
```

Just to be paranoid we "#uci show" to make sure UCI picks up the rules then we "#uci commit" and reboot the router.

at this point we have full access to the servers idrac6

### 5.45.3 Related Pages

---

#### **OpenVPN attempt #2**

[wiki:OpenVPNOnLEDE OpenVPN on LEDE]

#### **Adventures in deploying OpenWRT/LEDE**

- [wiki:OpenWRTonMR3020 Open WRT on TP-Link MR3020]
- [wiki:OpenWRTonLinkSysEA3500 Open WRT on LinkSYS EA3500]

## 5.46 LEDE EA3500 Note

---

This guy required me to upload the 15.04 and sys upgrade. Otherwise not a huge deal.

- <https://wiki.openwrt.org/toh/linksys/ea3500>



## 5.47 LEDE build (old)

---

### Building new firmware

```
root@sandbox:/home/openwrt/15.05.1/OpenWrt-ImageBuilder-15.05.1-ar71xx-generic.Linux-x86_64/bin/ar71xx# make info
....
root@sandbox:/home/openwrt/15.05.1/OpenWrt-ImageBuilder-15.05.1-ar71xx-generic.Linux-x86_64/bin/ar71xx# make image PROFILE=TLMR3020 PACKAGES="nano"
...
```

getting firmware onto local system. the stock firmware will not accept a firmware that is not the same name as a stock firmware.

```
viscious:vpn don$ scp feurig@sandbox:/home/openwrt/15.05.1/OpenWrt-ImageBuilder-15.05.1-ar71xx-generic.Linux-x86_64/bin/ar71xx/openwrt-15.05.1-ar71xx-generic-tl-mr3020-v1-squashfs-factory.bin .
viscious:vpn don$ mv openwrt-15.05.1-ar71xx-generic-tl-mr3020-v1-squashfs-factory.bin mr3020nv1_en_3_17_2_up_boot(150921).bin
```

At this point you can telnet to the router and reset the root password (which will disable telnet and enable ssh)

### related

- [wiki:LEDE LEDE]

### References

- <https://nicolas314.wordpress.com/2015/12/09/openwrt-on-mr3020/>
- <https://wolfgang.reutz.at/2012/04/12/openwrt-on-tp-link-mr3020-as-infopoint-with-local-webserver/>
- <https://blog.philippklaus.de/2012/03/openwrt-on-a-tp-link-tl-mr3020-router/>
- <https://openwrt.org/docs/guide-user/additional-software/imagebuilder>

## 5.48 LEDE E900 Notes

---

```
feurig@sandbox:~$ cd /home/openwrt/current/openwrt-imagebuilder-18.06.1-brcm47xx-mips74k.Linux-x86_64/
feurig@sandbox:/ho...64$ sudo cat ~joe/.ssh/authorized_keys ~feurig/.ssh/authorized_keys >files/etc/dropbear/authorized_keys
feurig@sandbox:/ho...64$ make image PROFILE=linksys-e900-v1 PACKAGES="nano sudo shadow shadow-utils shadow-vipw -luci -ppp -ppp-mod-pppoe -odhcp6c -odhcpd-ipv6only"
FILES="files/"
....
Calculating checksums...
feurig@sandbox:/ho...64$ ls bin/targets/brcm47xx/mips74k/
openwrt-brcm47xx-mips74k-asus-rt-ac53u-squashfs.trx
....
brcm47xx-mips74k-linksys-e900-v1-squashfs.bin
...
openwrt-brcm47xx-mips74k-linksys-e2500-v2.1-squashfs.bin ...
feurig@sandbox:/home/openwrt/current/openwrt-imagebuilder-18.06.1-brcm47xx-mips74k.Linux-x86_64$
```

## 5.49 LEDE Remote Syslog

Sending router system logs to remote server using rsyslog Either 18.04 Server or the LXC snap has rsyslog installed. So getting syslog information from the admin firewall is pretty simple. Its possible that we may need to provide a server other than kb2018 to make this ideal however I wanted to make sure that the syslogs stayed on the admin lan.

### 5.49.1 Sending logs to remote server

Modify the log configuration entries to point to the remote syslog and selecting a port and protocol is all that is needed.

```
feurig@knight:~$ cat /etc/config/system

config system
 option hostname 'knight'
 option timezone 'PDT'
 option ttylogin '0'
 option log_size '64'
 option urandom_seed '0'
 option log_ip '192.168.31.159'
 option log_port '514'
 option log_proto 'udp'

config timeserver 'ntp'
 option enabled '1'
 option enable_server '0'
 list server '0.lede.pool.ntp.org'
 list server '1.lede.pool.ntp.org'
 list server '2.lede.pool.ntp.org'
 list server '3.lede.pool.ntp.org'
```

Afterwards commit the configuration and restart the log daemon.

```
root@knight:/home/feurig# uci commit
root@knight:/home/feurig# /etc/init.d/log enable
root@knight:/home/feurig# /etc/init.d/log restart
```

### 5.49.2 Configuring rsyslogd on the remote server

Once you swim through the bagillian conflicting howtoo's for the multiple versions of rsyslogd you add the following lines to /etc/rsyslog.conf and restart it.

```
root@kb2018:/var/log# nano /etc/rsyslog.conf
....
provides UDP syslog reception
module(load="imudp")
input(type="imudp" port="514")

##Try exameple template for remote logs.
$template RemoteLogs, "/var/log/%HOSTNAME%/%PROGRAMNAME%.log"
```

. ?RemoteLogs .... root@kb2018:/var/log# service rsyslog restart

And test it.

```
root@knight:/home/feurig# logger testlog meh

root@kb2018:/var/log# tail /var/log/knight/
dropbear.log logread.log root.log sudo.log
root@kb2018:/var/log# tail /var/log/knight/dropbear.log
2018-12-21T18:50:54-08:00 knight dropbear[2465]: Exit (feurig): Keepalive timeout
2018-12-21T19:44:31-08:00 knight dropbear[2524]: Child connection from 193.193.70.69:59547
2018-12-21T19:44:31-08:00 knight dropbear[2524]: Exit before auth: Exited normally
2018-12-21T20:02:11-08:00 knight dropbear[2541]: Child connection from 111.43.34.166:2323
2018-12-21T20:02:12-08:00 knight dropbear[2541]: Exit before auth: Exited normally
2018-12-21T21:22:01-08:00 knight dropbear[2598]: Child connection from 35.159.6.209:37640
2018-12-21T21:22:13-08:00 knight dropbear[2598]: Login attempt for nonexistent user from 35.159.6.209:37640
2018-12-21T21:22:14-08:00 knight dropbear[2598]: Exit before auth: Disconnect received
2018-12-21T21:54:15-08:00 knight dropbear[2623]: Child connection from 97.115.132.190:59586
2018-12-21T21:54:17-08:00 knight dropbear[2623]: Pubkey auth succeeded for 'feurig' with key sha1!! 2a:26:75:a7:ec:fe:92:f4:b5:64:2e:26:26:dd:12:e5:d5:68:4f:67 from
97.115.132.190:59586
root@kb2018:/var/log# tail /var/log/knight/sudo.log
2018-12-21T18:35:42-08:00 knight sudo: feurig : TTY=pts/0 ; PwD=/home/feurig ; USER=root ; COMMAND=/bin/ash
2018-12-21T22:13:35-08:00 knight sudo: feurig : TTY=pts/0 ; PwD=/home/feurig ; USER=root ; COMMAND=/sbin/uci commit
root@kb2018:/var/log# tail /var/log/knight/root.log
2018-12-21T17:37:28-08:00 knight root: testLog "Blah1"
```

```
2018-12-21T18:35:54-08:00 knight root: testlog meh
root@kb2018:/var/log#
```

### Link Dump

- <https://forum.archive.openwrt.org/viewtopic.php?id=11912>
- <https://kuther.net/howtos/howto-log-firewall-openwrt-remote-rsyslog>
- <https://feeding.cloud.geek.nz/posts/debugging-openwrt-routers-by-shipping/>
- <https://www.rsyslog.com/storing-messages-from-a-remote-system-into-a-specific-file/>

## 5.50 Lets Encrypt Certificates

Recently a recruiter was unable to see my blog because the certificate was self signed. So I fixed the certificate on the blog and servrdocs with one of the EFF sponsored certs from LetsEncrypt. You can install certbot using the snap recommended by their docs or you can just

```
apt-get install certbot
```

Once its done you need to open 2 terminals on the target machine. In the first one you need to make whatever adjustments to your server to serve the url !<http://myserver.mydomain.whatever/.well-known/acme-challenge/xxxx> At which point you can run certbot and create your cert.

```
root@herbert:/var/www/html# certbot certonly --manual
Saving debug log to /var/log/letsencrypt/letsencrypt.log
Plugins selected: Authenticator manual, Installer None
Enter email address (used for urgent renewal and security notices) (Enter 'c' to
cancel): don@suspectdevices.com

Please read the Terms of Service at
https://letsencrypt.org/documents/LE-SA-v1.2-November-15-2017.pdf. You must
agree in order to register with the ACME server at
https://acme-v02.api.letsencrypt.org/directory

(A)gree/(C)ancel: A

Would you be willing to share your email address with the Electronic Frontier
Foundation, a founding partner of the Let's Encrypt project and the non-profit
organization that develops Certbot? We'd like to send you email about our work
encrypting the web, EFF news, campaigns, and ways to support digital freedom.

(Y)es/(N)o: Y
Please enter in your domain name(s) (comma and/or space separated) (Enter 'c'
to cancel): serverdocs.suspectdevices.com
Obtaining a new certificate
Performing the following challenges:
http-01 challenge for serverdocs.suspectdevices.com

NOTE: The IP of this machine will be publicly logged as having requested this
certificate. If you're running certbot in manual mode on a machine that is not
your server, please ensure you're okay with that.

Are you OK with your IP being logged?

(Y)es/(N)o: Y

Create a file containing just this data:

_bbuZOGf0JH0qF1F1LEAGe9s-e9b3IUyq6C1UUAg7xA.bvfoagN5gvQVzT-7dZuyvNibIYAUGx3MBNp0YLFo_g

And make it available on your web server at this URL:

http://serverdocs.suspectdevices.com/.well-known/acme-challenge/_bbuZOGf0JH0qF1F1LEAGe9s-e9b3IUyq6C1UUAg7xA

Press Enter to Continue
Waiting for verification...
Cleaning up challenges

IMPORTANT NOTES:
- Congratulations! Your certificate and chain have been saved at:
 /etc/letsencrypt/live/serverdocs.suspectdevices.com/fullchain.pem
 Your key file has been saved at:
 /etc/letsencrypt/live/serverdocs.suspectdevices.com/privkey.pem
 Your cert will expire on 2021-02-01. To obtain a new or tweaked
 version of this certificate in the future, simply run certbot
 again. To non-interactively renew *all* of your certificates, run
 "certbot renew"
- Your account credentials have been saved in your Certbot
 configuration directory at /etc/letsencrypt. You should make a
 secure backup of this folder now. This configuration directory will
 also contain certificates and private keys obtained by Certbot so
 making regular backups of this folder is ideal.
- If you like Certbot, please consider supporting our work by:

 Donating to ISRG / Let's Encrypt: https://letsencrypt.org/donate
 Donating to EFF: https://eff.org/donate-le

- We were unable to subscribe you the EFF mailing list because your
```

e-mail address appears to be invalid. You can try again later by visiting <https://act.eff.org>.

- <https://www.ssllabs.com/ssltest/analyze.html?d=blog.suspectdevices.com>

## 5.51 MigrateUsers

---

```
Migrate Users UID/GID
apt-get update
apt-get dist-upgrade
tasksel
apt-get update
apt-get dist-upgrade
vipw
export EDITOR=nano

chown --from=1000:1000 999:999 /. -Rv
```

## 5.52 Migrating Services to LXD

Up until 31 Jan 2019 medea was still providing critical services to the network and to myself. None of these services are disentangled enough to move them quickly. Starting with the web/mail servers we first attempted to set up a container on Medea and Migrate that container to bs2020. Adding a bridge to a running server with 30 aliases wasn't exactly straightforward so the services are being built on containers on bs2020 and migrated, Starting with trac.

### 5.52.1 osx-avr, suspectdevices.com, 3dangst, dns servers

#### track server Apache, postgres, trac. (trac.suspededevices.com/198.202.31.221)

This server could have been better documented but I needed it her to document everything else.

##### INSTALL NOTES

- Backed up old server according to <https://trac.edgewall.org/wiki/TracBackup#RestoringaBackup>
- installed everything from debian packages except for the wikiprint module which had to be manually installed.
- Moved trac to /var/www/trac (default document root was /var/www/html may move it again).
- path is hardcoded in cgi-bin/trac.wsgi
- The database file from hotcopy did not assign the database and tables to the trac\_db\_admin user. (manually fixed)
- .egg-cache and plugins directories must be owned by www-data
- replaced index.html with a redirect to /trac.
- created dns entry for trac.suspectdevices.com
- replaced apacheconfig on old server with Redirect

Redirect permanent /project/todo http://trac.suspectdevices.com/trac

#### Suspect devices wordpress blog

- create lxc container and install lamp server using tasksel.

```
root@bs2020:~# lxc init local:ubuntu12.04 ian -p susdev Creating ian root@bs2020:~# lxc start ian root@bs2020:~# lxc exec ian bash ... edit interfaces file and reboot or restart network services ... root@ian:~# apt-get install tasksel root@ian:~# tasksel ... select lamp server ... set password for mysql server ...
```

- Sort out the wordpress blog from the other legacy stuff.

```
root@medea:/home/newcourse/suspectdevices/www# ls -ls total 9916 4 drwxr-xr-x 4 www-data www-data 4096 Nov 17 2015 art2013 4 drwxr-xr-x 6 www-data www-data 4096 Jan 13 09:42 blahg 4 drwxrwxr-x 2 www-data staff 4096 Oct 10 2011 blog 5240 -rw-r--r-- 1 www-data root 5365300 Jun 22 2012 cma.tgz 4 drwxr-xr-x 3 www-data www-data 4096 Aug 25 2012 CookingWithMapleBacon 4 drwxrwxr-x 2 www-data staff 4096 Jan 14 2012 css 4 drwxrwxr-x 2 www-data staff 4096 Mar 1 2012 data 4 drwxrwxr-x 2 www-data staff 4096 Feb 12 2013 demo 4 -rw-rw-r-- 1 www-data staff 897 Nov 12 2011 dorkboard_gallery.html 8 -rw-rw-r-- 1 www-data staff 4890 Jan 16 2012 dorkboard.html 4 drwxrwxrwx 2 www-data staff 4096 Jun 30 2013 drop 4 -rw-rw-r-- 1 www-data staff 2970 Jun 22 2012 duce.html 0 -rw-rw-r-- 1 www-data staff 0 Nov 12 2011 favicon.ico 4 drwxr-xr-x 2 www-data www-data 4096 Feb 11 2013 feedme 4 drwxrwxr-x 3 www-data staff 4096 Nov 12 2011 images 4 -rw-r--r-- 1 www-data root 76 Jun 27 2012 index.php 4 drwxrwxr-x 3 www-data staff 4096 Nov 12 2011 js 4432 -rw-r--r-- 1 www-data root 4538093 Jun 22 2012 latest.tar.gz 4 drwxr-xr-x 2 www-data camo 4096 Jul 28 2012 library 4 -rw-rw-r-- 1 www-data staff 819 Nov 12 2011 others.html 4 drwxr-xr-x 19 www-data don 4096 Nov 4 2014 PCFA 4 -rw-rw-r-- 1 www-data staff 923 Oct 13 2011 pindex.php 4 drwxr-xr-x 2 www-data don 4096 Dec 9 2016 reference 4 drwxr-xr-x 2 www-data root 4096 Apr 8 2013 resumes 4 -rw-rw-r-- 1 www-data staff 1371 Mar 13 2017 static.html 4 -rw-rw-r-- 1 www-data staff 2599 Feb 26 2012 tad.html 4 drwxr-xr-x 3 www-data don 4096 Dec 19 2012 talks 28 -rw-rw-r-- 1 www-data staff 27241 Jun 22 2011 temp_bg.png 68 -rw-rw-r-- 1 www-data staff 68019 Jun 22 2011 temp_board.png 40 -rw-rw-r-- 1 www-data staff 38110 Jun 22 2011 temp_logo.png 4 drwxr-xr-x 3 www-data www-data 4096 Jun 27 2012 TheBaco-matic5000-OSB 0 lrwxrwxrwx 1 www-data root 5 Feb 15 2013 wordpress -> blahg 4 -rw-rw-r-- 1 www-data staff 3559 May 12 2012 workshops.html.old root@medea:/home/newcourse/suspectdevices/www# mkdir ../exodus root@medea:/home/newcourse/suspectdevices/www# cp -p *.html ../exodus root@medea:/home/newcourse/suspectdevices/www# cp -rpv talks/EpicMidiFail/
```



```
../exodus ... root@medea:/home/newcourse/suspectdevices/www# cp -rpv images ../exodus/ ... root@medea:/home/newcourse/
suspectdevices/www# cp -rpv blahg ../exodus/ ...
```

- dump the database

```
root@medea:/home/newcourse/suspectdevices/www# mysqldump -u www-data -p susdevweb> ../exodus/susdevweb.dump
Enter password:
```

- move and untar into /var/www/html

- restore database

```
root@ian:/var/www# mysqladmin -p create susdevweb Enter password: root@ian:/var/www/html/blahg# mysql -p
susdevweb< exodus/susdevweb.dump Enter password: root@ian:/var/www/html/blahg# mysql -p susdevweb Enter
password: ... mysql> CREATE USER 'www-data'@'localhost' IDENTIFIED BY 'somepassword'; Query OK, 0 rows affected
(0.00 sec) mysql> GRANT ALL PRIVILEGES ON *.* TO 'www-data'@'localhost'; Query OK, 0 rows affected (0.00 sec)

mysql>
```

- adjust /etc/apache2/sites-enabled/000-default

```
.... not really needed
```

- enable mod rewrite and .htaccess override.

```
root@ian:~# nano /etc/apache2/apache2.conf ... Options Indexes FollowSymLinks AllowOverride All Require all granted ...
root@ian:~# cd /etc/apache2/mods-enabled/ root@ian:/etc/apache2/mods-enabled# ln -s ../mods-available/rewrite.load .
root@ian:/etc/apache2/mods-enabled# apachectl configtest Syntax OK root@ian:/etc/apache2/mods-enabled# apachectl
restart
```

- route / to /blahg/ and check rewrite rules for wordpress site

```
root@ian:~# nano /var/www/html/.htaccess RewriteEngine on RewriteRule "^/$" "/blahg/" [R]

root@ian:~# cat /var/www/html/blahg/.htaccess RewriteEngine On RewriteBase /blahg/ RewriteRule ^index.php$ - [L]
RewriteCond %{REQUEST_FILENAME} !-f RewriteCond %{REQUEST_FILENAME} !-d RewriteRule . /blahg/index.php [L]
```

## 5.52.2 Static web server.

### busholini, Straight.fromhell.com, (with processing) osxavr.org

In order to mitigate the issues around CMS's such as wordpress, web sites whos primary purpose is to present photos and information that do not require dynamic content will be moved to a lighttpd server using named virtual hosts. Once this is tested it will be moved to 198.202.31.230 (formally www.suspectdevices.com)

- create lts container and apt-get install lighttpd
- copy static content into directories under /var/www
- edit /etc/lighttpd/lighttpd.conf

....

## 5.53 default server and configuration

```
server.document-root = "/var/www/busholini/www" server.upload-dirs = ("/var/cache/lighttpd/uploads") server.errorlog = "/var/log/lighttpd/error.log" server.pid-file = "/var/run/lighttpd.pid" server.username = "www-data" server.groupname = "www-data" server.port = 80
```

## 5.54

## 5.55 virtualhosts

## 5.56

```
$HTTP["host"] =~ "www.suspectdevices.com" { url.redirect # ("^/(.*)"> "http://blog.suspectdevices.com/$1") }
$HTTP["host"] =~ "(^|.)digithink.com$" { server.document-root = "/var/www/digithink/www" } $HTTP["host"] =~ "(^|.)thesofttargets.com$" { server.document-root = "/var/www/thesofttargets/www" }
```

## 5.57 disable php

```
index-file.names = ("index.html", "index.lighttpd.html") url.access-deny = ("~", ".inc", ".php")
```

Note/todo: the redirects should be more specific \* ie /project/todo -> trac.suspectdevices.com \* ie /blahg/ -> blog.suspectdevices.com

### 5.57.1 DNS/MAIL server (naomi)

#### DNS

- consolidate active zone files and create single master.conf to be included by /etc/bind/named.conf.local
- ```
// // Do any local configuration here //

// Consider adding the 1918 zones here, if they are not used in your // organization include "/etc/bind/zones/master.conf";
root@naomi:~# cat /etc/bind/zones/master.conf zone "digithink.com" in { type master; file "/etc/bind/zones/digithink.hosts"; };
zone "fromhell.com" in { type master; file "/etc/bind/zones/fromhell.hosts"; };
zone "busholini.org" in { type master; file "/etc/bind/zones/busholini.hosts"; };
zone "3dangst.com" in { type master; file "/etc/bind/zones/3dangst.hosts"; };
zone "osx-avr.org" in { type master; file "/etc/bind/zones/osx-avr.hosts"; };
zone "suspectdevices.com" { type master; file "/etc/bind/zones/suspectdevices.hosts"; };
zone "thesofttargets.com" { type master; file "/etc/bind/zones/thesofttargets.hosts"; };
```

```
zone "bresgal.com" in { type master; file "/etc/bind/zones/bresgal.hosts"; };
zone "bresgal.org" in { type master; file "/etc/bind/zones/bresgal.hosts"; };
zone "bluegin.net" in { type master; file "/etc/bind/zones/bluegin.hosts"; };
```

- check and restart bind

```
root@naomi:~# named-checkconf /etc/bind/named.conf root@naomi:~# named-checkconf /etc/bind/named.conf
root@naomi:~# service bind9 restart root@naomi:~# service bind9 status ● bind9.service - BIND Domain Name Server
Loaded: loaded (/lib/systemd/system/bind9.service; enabled; vendor preset: enabled) Drop-In: /run/systemd/generator/
bind9.service.d └─50-insserv.conf$named.conf Active: active (running) since Tue 2018-01-30 10:19:15 PST; 6s ago Docs:
man:named(8) Process: 962 ExecStop=/usr/sbin/rndc stop (code=exited, status=0/SUCCESS) Main PID: 965 (named)
CGroup: /system.slice/bind9.service └─965 /usr/sbin/named -f -u bind
```

```
Jan 30 10:19:15 naomi named[965]: zone bresgal.org/IN: sending notifies (serial 2009123000) Jan 30 10:19:15 naomi
named[965]: zone suspectdevices.com/IN: sending notifies (serial 2018012902) Jan 30 10:19:15 naomi named[965]: zone
3dangst.com/IN: sending notifies (serial 2004072801) Jan 30 10:19:15 naomi named[965]: zone busholini.org/IN: sending
notifies (serial 2018012201) Jan 30 10:19:15 naomi named[965]: zone osx-avr.org/IN: sending notifies (serial 2005032100)
Jan 30 10:19:15 naomi named[965]: zone digithink.com/IN: sending notifies (serial 2018012200) Jan 30 10:19:15 naomi
named[965]: zone fromhell.com/IN: sending notifies (serial 2004072000) Jan 30 10:19:15 naomi named[965]: zone
bluegin.net/IN: sending notifies (serial 2004072500) Jan 30 10:19:15 naomi named[965]: zone thesofttargets.com/IN:
sending notifies (serial 2018012200) Jan 30 10:19:15 naomi named[965]: zone bresgal.com/IN: sending notifies (serial
2009123000)
```

- install bind9 and email services via tasksel
- move dns1 ip from medea to naomi
- reboot servers.

Mail

Based on the file dates of the Maildir's being updated by postfix on the old server..

- Look at existing server for active email users.

```
root@medea:~# find / -name Maildir -a -newer www/postgres7JUL17.dump -print /var/www/Maildir /home/eldufe/Maildir /
home/don/Maildir /home/fromhell/users/feurig/Maildir
```

We notice that only three users are reading email so we need to serve those users.

- So create users for feurig@fromhell.com, eldufe@busholini.org and don@suspectdevices.com since www is going to be exclusively spam.

```
root@naomi:~# useradd -c "The Commander and Thief" -m eldufe root@naomi:~# useradd -c "D Delmar Davis" -m don
```

The rest of the documentation has been moved to a separate [wiki:UbuntuMailServerSetup mail server setup] document.

Secondary DNS Server

- create server
- install dns using tasksel
- transfer and convert master configuration to slave.

```
root@teddy:~# cd /etc/bind root@teddy:/etc/bind# mkdir zones root@teddy:/etc/bind# scp don@198.202.31.231:/etc/bind/
zones/master.conf slave.conf The authenticity of host '198.202.31.231 (198.202.31.231)' can't be established. ECDSA key
fingerprint is SHA256:WFKs+2xinTQKgPhIM6fjCy2FMpY4SbeYvM2lQZpifi. Are you sure you want to continue connecting
(yes/no)? yes Warning: Permanently added '198.202.31.231' (ECDSA) to the list of known hosts. don@198.202.31.231's
password: master.conf 100% 983 1.0KB/s 00:00
root@teddy:/etc/bind# sed 's/master;/slave;'\n\tnmasters { 198.202.31.141; };/' slave.conf >zones/slave.conf root@teddy:/etc/
bind# nano named.conf.local

// Do any local configuration here //
```

```
// Consider adding the 1918 zones here, if they are not used in your // organization //include "/etc/bind/zones.rfc1918";
include "/etc/bind/zones/slave.conf";
```

- deal with duplicate filename and slave configuration in bresgals....

```
root@teddy:/etc/bind# named-checkconf /etc/bind/zones/slave.conf:52: writeable file '/etc/bind/zones/bresgal.hosts': already
in use: /etc/bind/zones/slave.conf:46 root@teddy:/etc/bind# nano /etc/bind/zones/slave.conf .... root@teddy:/etc/bind# service
bind9 restart root@teddy:/etc/bind# service bind9 status ● bind9.service - BIND Domain Name Server Loaded: loaded (/lib/
systemd/systemd/bind9.service; enabled; vendor preset: enabled) Drop-In: /run/systemd/generator/bind9.service.d └─50-
insserv.conf-$named.conf Active: active (running) since Wed 2018-01-31 22:17:00 PST; 5min ago Docs: man:named(8)
Process: 5436 ExecStop=/usr/sbin/rndc stop (code=exited, status=1/FAILURE) Main PID: 5450 (named) Tasks: 27 Memory:
30.4M CPU: 114ms CGroup: /system.slice/bind9.service └─5450 /usr/sbin/named -f -u bind
```

```
Jan 31 22:17:01 teddy named[5450]: zone bluegin.net/IN: transferred serial 2004072500 Jan 31 22:17:01 teddy
named[5450]: transfer of 'bluegin.net/IN' from 198.202.31.141#53: Transfer status: success Jan 31 22:17:01 teddy
named[5450]: transfer of 'bluegin.net/IN' from 198.202.31.141#53: Transfer completed: 1 messages, Jan 31 22:17:01 teddy
named[5450]: zone bresgal.org/IN: transferred serial 2009123000 Jan 31 22:17:01 teddy named[5450]: zone bluegin.net/
IN: sending notifies (serial 2004072500) Jan 31 22:17:01 teddy named[5450]: transfer of 'bresgal.org/IN' from
198.202.31.141#53: Transfer status: success Jan 31 22:17:01 teddy named[5450]: transfer of 'bresgal.org/IN' from
198.202.31.141#53: Transfer completed: 1 messages, Jan 31 22:17:01 teddy named[5450]: zone bresgal.org/IN: sending
notifies (serial 2009123000) Jan 31 22:17:01 teddy named[5450]: dumping master file: /etc/bind/zones/tmp-qGurg6XtTG:
open: permission denied Jan 31 22:17:01 teddy named[5450]: dumping master file: /etc/bind/zones/tmp-jUyE6xKRDk: open:
permission denied
```

- Move zone files to /var/lib/bind/ because apparmor won't let you write to /etc/bind/zones...

```
root@teddy:~# sed -i 's/etc/bind/zones/var/lib/bind/' /etc/bind/zones/slave.conf root@teddy:~# service bind9 restart
root@teddy:~# tail /var/log/syslog Sep 8 13:48:56 teddy named[7118]: zone bresgal.com/IN: sending notifies (serial
2009123000) Sep 8 13:48:56 teddy named[7118]: transfer of 'bluegin.net/IN' from 198.202.31.141#53: connected using
198.202.31.132#45499 Sep 8 13:48:56 teddy named[7118]: zone suspectdevices.com/IN: transferred serial 2018080300
Sep 8 13:48:56 teddy named[7118]: transfer of 'suspectdevices.com/IN' from 198.202.31.141#53: Transfer status: success
Sep 8 13:48:56 teddy named[7118]: transfer of 'suspectdevices.com/IN' from 198.202.31.141#53: Transfer completed: 1
messages, 32 records, 1228 bytes, 0.001 secs (1228000 bytes/sec) Sep 8 13:48:56 teddy named[7118]: zone
suspectdevices.com/IN: sending notifies (serial 2018080300) Sep 8 13:48:56 teddy named[7118]: zone bluegin.net/IN:
transferred serial 2004072500 Sep 8 13:48:56 teddy named[7118]: transfer of 'bluegin.net/IN' from 198.202.31.141#53:
Transfer status: success Sep 8 13:48:56 teddy named[7118]: transfer of 'bluegin.net/IN' from 198.202.31.141#53: Transfer
completed: 1 messages, 18 records, 450 bytes, 0.001 secs (450000 bytes/sec) Sep 8 13:48:56 teddy named[7118]: zone
bluegin.net/IN: sending notifies (serial 2004072500) root@teddy:~# ls /var/lib/bind/ 3dangst.hosts bluegin.hosts
bresgal1.hosts digithink.hosts osx-avr.hosts thesofttargets.hosts bind9-default.md5sum bresgal0.hosts busholini.hosts
fromhell.hosts suspectdevices.hosts root@teddy:~#
```

5.57.2 Sidenote: 17.10/18.04 container

While we were running up new containers we started the process of looking at the changes coming down the road (next LTS candidate) [BleedingEdgeServer Phillip] is our current exploration into what the kids are up to.

- BleedingEdgeServer

5.57.3 Linkdump

- <https://stackoverflow.com/questions/33377916/migrating-lxc-to-lxd>
- <https://bobcares.com/blog/wordpress-hosting-using-lxd-lxc-server-virtualization-solution/3/>
- <https://wparena.com/how-to-move-a-wordpress-site-from-one-server-to-another/>
- <https://www.quora.com/How-do-you-export-a-WordPress-site-to-a-static-HTML-i-e-how-do-you-remove-all-WordPress-functionality-from-a-WordPress-theme-to-turn-it-into-a-plain-HTML-theme-and-are-there-any-%E2%80%98export-as-HTML%E2%80%99-type-features-available>
- <https://stackoverflow.com/questions/17468109/postfix-unable-to-find-etc-postfix-virtual-file>
- <https://wordpress.org/plugins/static/>

- <https://wordpress.org/plugins/static-html-output-plugin/>
- <https://zargony.com/2008/02/04/migrating-from-apache-to-lighttpd-with-name-based-virtual-hosts-and-ssl/>
- <https://help.ubuntu.com/community/MailServer>
- <https://help.ubuntu.com/community/Dovecot>
- <https://help.ubuntu.com/community/Postfix>
- <https://help.ubuntu.com/lts/serverguide/postfix.html>
- <https://linuxide.com/ubuntu-how-to/setup-postfix-dovecot-mysql-ubuntu-1604/>
- <https://www.tecmint.com/setup-postfix-mail-server-in-ubuntu-debian/>
- <https://www.linuxbabe.com/mail-server/secure-email-server-ubuntu-16-04-postfix-dovecot>
- <https://skrilnetz.net/setup-your-own-mailserver/>
- <https://askubuntu.com/questions/54960/how-do-i-set-up-an-email-server#55027>
- <https://www.digitalocean.com/community/tutorials/how-to-install-and-configure-postfix-on-ubuntu-16-04>
- http://www.postfix.org/COMPATIBILITY_README.html
- <https://unix.stackexchange.com/questions/145771/mail-filtering-with-procmail-in-a-postfix-dovecot-system-with-virtual-users>
- <https://www.exratione.com/2016/05/a-mailserver-on-ubuntu-16-04-postfix-dovecot-mysql/>
- <https://www.digitalocean.com/community/tutorials/how-to-install-and-configure-postfix-as-a-send-only-smtp-server-on-ubuntu-16-04>
- http://www.postfix.org/STANDARD_CONFIGURATION_README.html#null_client
- <https://askubuntu.com/questions/967091/zpool-degrades-when-plugging-in-a-drive>

5.58 Mullein

5.58.1 Update Me

Mullein is an Asus mips based router which was initially set up as a [wiki:OpenVPNOnLEDE VPN server] for the house providing an internal lan for personal data and IOT Projects

5.59 New Trac Container

5.59.1 background

Our trac server has been setup using the old trac implementation running on 16.04. It works but needs to be updated and cleaned up.

PHILOSOPHY

- Where possible use only ubuntu/debian supported packages as apposed to manual/pip so that updates can be kept abreast of.
- Move the excellent online documentation to a separate section so that copies (pdf books, static html, etc) of the site include only the relevant pages.
- Leverage lxc container to create a reusable trac image.
- Leverage lxc container to create a backup of the old content.
- Add SSL functionality.

LINKDUMP / REFERENCES

- <https://trac.edgewall.org/wiki/TracInstall>
- <https://trac.edgewall.org/wiki/TracModWSGI#ConfiguringAuthentication>
- <https://www.hiroom2.com/2018/11/16/ubuntu-1810-trac-en/>
- <https://seattle.poly.edu/wiki/TracModWSGI>
- https://github.com/viktorTarasov/OpenSC-SM/wiki/Trac-and-mod_wsgi
- <https://help.ubuntu.com/community/TracApacheModWsgi>
- https://blog.niklasottosson.com/linux/setup-trac-project-on-debian-wheezy-with-apache-using-the-mod_wsgi-and-basic-authentication/
- <https://stackoverflow.com/questions/6097515/deleting-trac-tickets-created-since-a-certain-date-until-today>

RAW DUMP OF INSTALL

```
root@douglas:~# apt-get install git
...
root@douglas:~# apt-get install mercurial
...
root@douglas:~# apt-get install postgresql
...
root@douglas:~# apt-get install python-psycopg2
...
root@douglas:~# apt-get install trac
...
Suggested packages:
  www-browser apache2-doc apache2-suexec-pristine | apache2-suexec-custom libjs-jquery-ui-docs liblcms2-utils fonts-linuxlibertine | ttf-linux-libertine texlive-lang-french
  texlive-latex-base texlive-latex-recommended doc-base python-genshi-doc python-pil-doc python-pil-dbg ttf-bitstream-vera python-setuptools-doc python-subversion-dbg
  sgml-base-doc libapache2-mod-wsgi python-textile trac-accountmanager trac-authopenid trac-bitten trac-bzr trac-customfieldadmin trac-email2trac trac-graphviz trac-javascript
  trac-mastertickets trac-mercurial trac-spamfilter trac-wikiprint trac-wikirename trac-wysiwyg trac-xmlrpc debhelper
...
root@douglas:~# apt-get update&&apt-get dist-upgrade&&apt-get auto remove
... go back to his t.. lxc file push douglas /usr/local/bin/update.sh ....
root@douglas:~# chmod 774 /usr/local/bin/update.sh
root@douglas:~# update.sh
----- begin updating douglas -----
...
=====### done=====
root@douglas:~# apt-cache search trac
... way too much crap here ...
python-offtrac - Python-based xmlrpc client library for trac instances (Python 2)
....
trac - Enhanced wiki and issue tracking system for software development projects
trac-accountmanager - account management plugin for Trac
trac-announcer - enhanced e-mail notification system for Trac
trac-authopenid - OpenID authentication plugin for Trac
trac-bitten - continuous integration plugin for Trac
trac-bitten-slave - continuous integration plugin for Trac
trac-codecomments - code comments and review plugin for Trac
trac-customfieldadmin - panel for administrating custom ticket fields in Trac
trac-datefield - Add custom date fields to Trac tickets
trac-diaxiview - Renders dia and vdx files in Trac
trac-email2trac - Creates and amends Trac tickets from e-mail
```

```

trac-graphviz - Graphs printing plugin for Trac
trac-httpauth - Force HTTP authentication from within Trac
trac-icalview - Provides iCalendar feeds for ticket queries
trac-includemacro - Include external resources in a Trac wiki page
trac-jsganttt - displays Trac tickets as a Gantt chart in a wiki page
trac-mastertickets - adds inter-ticket dependencies to Trac
trac-mercurial - Mercurial version control backend for Trac
trac-navadd - add custom items to main and meta navigation bar in Trac webapp
trac-privatetickets - Allows Trac users to only see tickets they are associated with
trac-privateticketsplugin - transitional dummy package for trac-privatetickets
trac-privatewiki - add private wiki ability to Trac
trac-roadmap - enhances the Trac roadmap with sorting and filtering
trac-sensitivetickets - Plugin for Trac ticketing system to hide tickets marked as sensitive
trac-spamfilter - Spam-prevention plugin for Trac
trac-subcomponents - use multiple layers of components in Trac
trac-subtickets - sub-ticket feature for Trac tickets
trac-tags - Tagging plugin for Trac wiki and issue tracking system
trac-translatedpages - Show translated versions of wiki page in the Trac web application
trac-virtualticketpermissions - Extended permissions plugin for Trac ticketing system
trac-wikiprint - Make Trac wiki pages printable, exporting to PDF or printable HTML
trac-wikitablemacro - SQL Table in Wiki Page for Trac
trac-wysiwyg - WYSIWYG style editor for the Trac issue tracking system
trac-xmlrpc - XML-RPC interface to the Trac wiki and issue tracking system
...
root@douglas:~# apt-cache search psycpg*
python-psycpg2 - Python module for PostgreSQL
python-psycpg2-dbg - Python module for PostgreSQL (debug extension)
python-psycpg2-doc - Python module for PostgreSQL (documentation package)
python3-psycpg2 - Python 3 module for PostgreSQL
python3-psycpg2-dbg - Python 3 module for PostgreSQL (debug extension)
python-psycpgreen - psycpg2 integration with coroutine libraries
python3-aiopg - PostgreSQL integration with asyncio
root@douglas:~# apt-get install python3-psycpg2
Reading package lists... Done
Building dependency tree
Reading state information... Done
Suggested packages:
  python-psycpg2-doc
The following NEW packages will be installed:
  python3-psycpg2
0 upgraded, 1 newly installed, 0 to remove and 0 not upgraded.
Need to get 152 kB of archives.
After this operation, 838 kB of additional disk space will be used.
Get:1 http://archive.ubuntu.com/ubuntu bionic/main amd64 python3-psycpg2 amd64 2.7.4-1 [152 kB]
Fetched 152 kB in 1s (178 kB/s)
Selecting previously unselected package python3-psycpg2.
(Reading database ... 56846 files and directories currently installed.)
Preparing to unpack .../python3-psycpg2_2.7.4-1_amd64.deb ...
Unpacking python3-psycpg2 (2.7.4-1) ...
Setting up python3-psycpg2 (2.7.4-1) ...
root@douglas:~# apt-get install libapache2-mod-wsgi python-textile trac-accountmanager trac-authopenid trac-bitten trac-customfieldadmin trac-email2trac trac-graphviz
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  adwaita-icon-theme at-spi2-core dconf-gsettings-backend dconf-service fontconfig fontconfig-config fonts-dejavu-core fonts-liberation glib-networking glib-networking-
common
  glib-networking-services graphviz gsettings-desktop-schemas gtk-update-icon-cache hicolor-icon-theme humanity-icon-theme libann0 libatk-bridge2.0-0 libatk1.0-0
libatk1.0-data
  libatspi2.0-0 libavahi-client3 libavahi-common-data libavahi-common3 libcairo-gobject2 libcairo2 libcdt5 libcgraph6 libcolord2 libcroc3 libcups2 libdatrie1 libdconf1
libegl-mesa0 libegl1 libepoxy0 libfontconfig1 libgbm1 libgd3 libgdk-pixbuf2.0-0 libgdk-pixbuf2.0-bin libgdk-pixbuf2.0-common libglapi-mesa libglvnd0 libgraphite2-3
libgtk-3-0
  libgtk-3-bin libgtk-3-common libgts-0.7-5 libgts-bin libgvc6 libgvpr2 libharfbuzz0b libice6 libjs-flot libjs-jquery-flot libjson-glib-1.0-0 libjson-glib-1.0-common
liblab-gamut1 libltdl7 libpango-1.0-0 libpangocairo-1.0-0 libpangoft2-1.0-0 libpathplan4 libpixmap-1-0 libproxy1v5 libpython2.7 librest-0.7-0 librsvg-2 librsvg2-bin
librsvg2-common libsm6 libsoup-gnome2.4-1 libsoup2.4-1 libthai-data libthai0 libwayland-client0 libwayland-cursor0 libwayland-egl1-mesa libwayland-server0 libx11-xcb1
libxaw7
  libxcb-dri2-0 libxcb-dri3-0 libxcb-present0 libxcb-render0 libxcb-shm0 libxcb-sync1 libxcb-xfixes0 libxcomposite1 libxcursor1 libxdamage1 libxf86vm3 libxi6
libxinerama1
  libxkbcommon0 libxmu6 libxpm4 libxrandr2 libxrender1 libxshmfence1 libxt6 libxtst6 python-html5lib python-openid python-pygraphviz python-six python-webencodings
trac-bitten-slave ubuntu-mono x11-common
Suggested packages:
  gsfonts graphviz-doc colord cups-common libgd-tools gvfs libjs-jquery-flot-docs python-lxml python-pygraphviz-doc python-regex getmail4
The following NEW packages will be installed:
  adwaita-icon-theme at-spi2-core dconf-gsettings-backend dconf-service fontconfig fontconfig-config fonts-dejavu-core fonts-liberation glib-networking glib-networking-
common
  glib-networking-services graphviz gsettings-desktop-schemas gtk-update-icon-cache hicolor-icon-theme humanity-icon-theme libann0 libapache2-mod-wsgi libatk-bridge2.0-0
libatk1.0-0 libatk1.0-data libatspi2.0-0 libavahi-client3 libavahi-common-data libavahi-common3 libcairo-gobject2 libcairo2 libcdt5 libcgraph6 libcolord2 libcroc3
libcups2
  libdatrie1 libdconf1 libegl-mesa0 libegl1 libepoxy0 libfontconfig1 libgbm1 libgd3 libgdk-pixbuf2.0-0 libgdk-pixbuf2.0-bin libgdk-pixbuf2.0-common libglapi-mesa
libglvnd0
  libgraphite2-3 libgtk-3-0 libgtk-3-bin libgtk-3-common libgts-0.7-5 libgts-bin libgvc6 libgvpr2 libharfbuzz0b libice6 libjs-flot libjs-jquery-flot libjson-glib-1.0-0
libjson-glib-1.0-common liblab-gamut1 libltdl7 libpango-1.0-0 libpangocairo-1.0-0 libpangoft2-1.0-0 libpathplan4 libpixmap-1-0 libproxy1v5 libpython2.7 librest-0.7-0
librsvg-2 librsvg2-bin librsvg2-common libsm6 libsoup-gnome2.4-1 libsoup2.4-1 libthai-data libthai0 libwayland-client0 libwayland-cursor0 libwayland-egl1-mesa
libwayland-server0 libx11-xcb1 libxaw7 libxcb-dri2-0 libxcb-dri3-0 libxcb-present0 libxcb-render0 libxcb-shm0 libxcb-sync1 libxcb-xfixes0 libxcomposite1 libxcursor1
libxdamage1 libxf86vm3 libxi6 libxinerama1 libxkbcommon0 libxmu6 libxpm4 libxrandr2 libxrender1 libxshmfence1 libxt6 libxtst6 python-html5lib python-openid python-
pygraphviz
  python-six python-textile python-webencodings trac-accountmanager trac-authopenid trac-bitten trac-bitten-slave trac-customfieldadmin trac-email2trac trac-graphviz
ubuntu-mono
  x11-common
0 upgraded, 119 newly installed, 0 to remove and 79 not upgraded.
Need to get 18.0 MB of archives.
After this operation, 83.3 MB of additional disk space will be used.
Do you want to continue? [Y/n]

```



```

root@douglas:~# apt-get install trac-mastertickets trac-mercurial trac-spamfilter trac-wikiprint trac-xmlrpc debhelper
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  autoconf automake autopoint autotools-dev binutils binutils-common binutils-x86-64-linux-gnu build-essential cpp cpp-7 dh-autoreconf dh-strip-nondeterminism dpkg-dev
  fakeroot
  g++ g++-7 gcc gcc-7 gcc-7-base gcc-8-base gettext gsfnts intltool-debian libalgorithm-diff-perl libalgorithm-diff-xs-perl libalgorithm-merge-perl libarchive-cpio-perl
  libarchive-zip-perl libart-2.0-2 libasan4 libatomic1 libbinutils libc-dev-bin libc6-dev libcc1-0 libcilkrts5 libdpkg-perl libfakeroot libfile-fcntllock-perl
  libfile-stripnondeterminism-perl libgcc-7-dev libgcc1 libgomp1 libisl19 libitm1 liblsan0 libltdl-dev libmail-sendmail-perl libmpc3 libmpx2 libquadmath0 libstdc++-7-dev
  libstdc++6 libsys-hostname-long-perl libtime-date-perl libtool libtsan0 libubsan0 linux-libc-dev m4 make manpages-dev mercurial mercurial-common po-debconf python-dns
  python-dnspython python-httplib2 python-lockfile python-pypdf2 python-renderpm python-reportlab python-reportlab-accell python-xhtml2pdf spambayes
Suggested packages:
  autoconf-archive gnu-stardards autoconf-doc binutils-doc cpp-doc gcc-7-locales dh-make dwz debian-keyring g++-multilib g++-7-multilib gcc-7-doc libstdc++6-7-dbg gcc-
  multilib
  flex bison gdb gcc-doc gcc-7-multilib libgcc1-dbg libgomp1-dbg libitm1-dbg libatomic1-dbg libasan4-dbg liblsan0-dbg libtsan0-dbg libubsan0-dbg libcilkrts5-dbg libmpx2-
  dbg
  libquadmath0-dbg gettext-doc libasprintf-dev libgettextpo-dev glibc-doc bzip libtool-doc libstdc++-7-doc gfortran | fortran95-compiler gcj-jdk m4-doc make-doc kdiff3
  | kdiff3-qt | kompare | meld | tkcvs | mgedit qct python-mysqldb python-openssl wish libmail-box-perl python-lockfile-doc python-renderpm-dbg pdf-viewer
  python-egenix-mxtexttools python-reportlab-doc
The following NEW packages will be installed:
  autoconf automake autopoint autotools-dev binutils binutils-common binutils-x86-64-linux-gnu build-essential cpp cpp-7 debhelper dh-autoreconf dh-strip-nondeterminism
  dpkg-dev
  fakeroot g++ g++-7 gcc gcc-7 gcc-7-base gettext gsfnts intltool-debian libalgorithm-diff-perl libalgorithm-diff-xs-perl libalgorithm-merge-perl libarchive-cpio-perl
  libarchive-zip-perl libart-2.0-2 libasan4 libatomic1 libbinutils libc-dev-bin libc6-dev libcc1-0 libcilkrts5 libdpkg-perl libfakeroot libfile-fcntllock-perl
  libfile-stripnondeterminism-perl libgcc-7-dev libgcc1 libgomp1 libisl19 libitm1 liblsan0 libltdl-dev libmail-sendmail-perl libmpc3 libmpx2 libquadmath0 libstdc++-7-dev
  libsys-hostname-long-perl libtime-date-perl libtool libtsan0 libubsan0 linux-libc-dev m4 make manpages-dev mercurial mercurial-common po-debconf python-dns python-
  dnspython
  python-httplib2 python-lockfile python-pypdf2 python-renderpm python-reportlab python-reportlab-accell python-xhtml2pdf spambayes trac-mastertickets trac-mercurial
  trac-spamfilter trac-wikiprint trac-xmlrpc
The following packages will be upgraded:
  gcc-8-base libgcc1 libstdc++6
3 upgraded, 78 newly installed, 0 to remove and 76 not upgraded.
Need to get 48.6 MB of archives.
After this operation, 197 MB of additional disk space will be used.
Do you want to continue? [Y/n]
.....
root@douglas:~# nano /etc/postgresql/10/main/pg_hba.conf
root@douglas:~# su - postgres
postgres@douglas:~$ psql template1
psql (10.6 (Ubuntu 10.6-0ubuntu0.18.04.1))
Type "help" for help.

template1=# create database tracdb with encoding = 'utf8';
CREATE DATABASE
template1=# create user tracuser password 'password';
CREATE ROLE
template1=# grant all on database tracdb to tracuser;
GRANT
template1=# \q
postgres@douglas:~$ exit
logout
..... grumble grumble ..... bad password ....
root@douglas:~# service postgres reload
postgres: unrecognized service
root@douglas:~# service postgresql reload
root@douglas:~# mkdir /var/
backups/ cache/ crash/ lib/ local/ lock/ log/ mail/ opt/ run/ snap/ spool/ tmp/ www/
root@douglas:~# mkdir /var/trac/devel
mkdir: cannot create directory '/var/trac/devel': No such file or directory
root@douglas:~# mkdir /var/trac/
root@douglas:~# mkdir /var/trac/devel
root@douglas:~# cd /var/trac/devel/
root@douglas:/var/trac/devel# mkdir repo env
root@douglas:/var/trac/devel# trac-admin /var/trac/devel/env/ initenv
Creating a new Trac environment at /var/trac/devel/env

Trac will first ask a few questions about your environment
in order to initialize and prepare the project database.

Please enter the name of your project.
This name will be used in page titles and descriptions.

Project Name [My Project]> Development

Please specify the connection string for the database to use.
By default, a local SQLite database is created in the environment
directory. It is also possible to use an existing MySQL or
PostgreSQL database (check the Trac documentation for the exact
connection string syntax).

Database connection string [sqlite:db/trac.db]> postgres://tracuser:password@localhost/tracdb

Creating and Initializing Project
Installing default wiki pages
InterWiki imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/InterWiki
WikiProcessors imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiProcessors
TracUpgrade imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracUpgrade
TracUnicode imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracUnicode
WikiPageNames imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiPageNames
TracRevisionLog imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracRevisionLog
TracWiki imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracWiki
TracSearch imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracSearch
TracGuide imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracGuide

```

```

TracLinks imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracLinks
TracInterfaceCustomization imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracInterfaceCustomization
TracBrowser imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracBrowser
TracTickets imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracTickets
WikiNewPage imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiNewPage
TracSupport imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracSupport
TracStandalone imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracStandalone
TracChangeLog imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracChangeLog
TracNavigation imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracNavigation
TracAccessibility imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracAccessibility
TracSyntaxColoring imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracSyntaxColoring
TracFineGrainedPermissions imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracFineGrainedPermissions
TracInstall imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracInstall
InterTrac imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/InterTrac
WikiMacros imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiMacros
TracImport imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracImport
TitleIndex imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TitleIndex
SandBox imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/SandBox
TracCgi imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracCgi
TracBackup imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracBackup
WikiHtml imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiHtml
TracTicketsCustomFields imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracTicketsCustomFields
CamelCase imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/CamelCase
TracModWsgi imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracModWsgi
WikiFormatting imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiFormatting
RecentChanges imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/RecentChanges
TracBatchModify imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracBatchModify
TracRepositoryAdmin imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracRepositoryAdmin
InterMapTxt imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/InterMapTxt
TracRoadmap imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracRoadmap
WikiDeletePage imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiDeletePage
TracWorkflow imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracWorkflow
WikiRestructuredText imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiRestructuredText
TracIni imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracIni
TicketQuery imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TicketQuery
TracNotification imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracNotification
TracEnvironment imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracEnvironment
TracPlugins imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracPlugins
WikiStart imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiStart
TracReports imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracReports
TracAdmin imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracAdmin
WikiRestructuredTextLinks imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/WikiRestructuredTextLinks
TracChangeset imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracChangeset
TracQuery imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracQuery
TracFastCgi imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracFastCgi
TracRss imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracRss
TracTimeline imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracTimeline
TracModPython imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracModPython
TracLogging imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracLogging
PageTemplates imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/PageTemplates
TracPermissions imported from /usr/lib/python2.7/dist-packages/trac/wiki/default-pages/TracPermissions

```

Project environment for 'Development' created.

You may now configure the environment by editing the file:

```
/var/trac/devel/env/conf/trac.ini
```

If you'd like to take this new project environment for a test drive,
try running the Trac standalone web server `tracd`:

```
tracd --port 8000 /var/trac/devel/env
```

Then point your browser to <http://localhost:8000/env>.
There you can also browse the documentation for your installed
version of Trac, including information on further setup (such as
deploying Trac to a real web server).

The latest documentation can also always be found on the project
website:

```
http://trac.edgewall.org/
```

Congratulations!

```

root@douglas:/var/trac/devel# tracd --port 8000 /media/shared/Admin/trac/repo/
...
root@douglas:/var/trac/devel# trac-admin /var/trac/devel/env/ deploy /var/trac/devel/www/
Copying resources from:
  trac.web.chrome.Chrome
  /usr/lib/python2.7/dist-packages/trac/htdocs
  /var/trac/devel/env/htdocs
Creating scripts.
root@douglas:/var/trac/devel# nano /etc/apache2/sites-enabled/000-default.conf
root@douglas:/var/trac/devel# service apache2 reload
root@douglas:/var/trac/devel# chmod u+x www/cgi-bin/trac.wsgi
root@douglas:/var/trac/devel# chown www-data env/conf/trac.ini
root@douglas:/var/trac/devel# trac-admin
trac-admin - The Trac Administration Console 1.2

```

Usage: trac-admin </path/to/projenv> [command [subcommand] [option ...]]

Invoking trac-admin without command starts interactive mode.

```

help      Show documentation
initenv   Create and initialize a new environment
root@douglas:/var/trac/devel# trac-admin /var/trac/devel/env/
Welcome to trac-admin 1.2
Interactive Trac administration console.
Copyright (C) 2003-2013 Edgewall Software

Type: '?' or 'help' for help on commands.

Trac [/var/trac/devel/env]> ?
trac-admin - The Trac Administration Console 1.2
help      Show documentation
initenv   Create and initialize a new environment
attachment add      Attach a file to a resource
attachment export   Export an attachment from a resource to a file or stdout
attachment list     List attachments of a resource
attachment remove   Remove an attachment from a resource
changeset added     Notify trac about changesets added to a repository
changeset modified  Notify trac about changesets modified in a repository
component add       Add a new component
component chown     Change component ownership
component list      Show available components
component remove    Remove/uninstall a component
component rename    Rename a component
config get          Get the value of the given option in "trac.ini"
config remove      Remove the specified option from "trac.ini"
config set          Set the value for the given option in "trac.ini"
deploy             Extract static resources from Trac and all plugins
hotcopy            Make a hot backup copy of an environment
milestone add       Add milestone
milestone completed Set milestone complete date
milestone due       Set milestone due date
milestone list      Show milestones
milestone remove    Remove milestone
milestone rename    Rename milestone
permission add       Add a new permission rule
permission export    Export permission rules to a file or stdout as CSV
permission import    Import permission rules from a file or stdin as CSV
permission list      List permission rules
permission remove    Remove a permission rule
priority add         Add a priority value option
priority change      Change a priority value
priority list        Show possible ticket priorities
priority order       Move a priority value up or down in the list
priority remove      Remove a priority value
repository add       Add a source repository
repository alias     Create an alias for a repository
repository list      List source repositories
repository remove    Remove a source repository
repository resync    Re-synchronize trac with repositories
repository set       Set an attribute of a repository
repository sync      Resume synchronization of repositories
resolution add       Add a resolution value option
resolution change    Change a resolution value
resolution list      Show possible ticket resolutions
resolution order     Move a resolution value up or down in the list
resolution remove    Remove a resolution value
session add          Create a session for the given sid
session delete       Delete the session of the specified sid
session list         List the name and email for the given sids
session purge        Purge anonymous sessions older than the given age or date
session set          Set the name or email attribute of the given sid
severity add         Add a severity value option
severity change      Change a severity value
severity list        Show possible ticket severities
severity order       Move a severity value up or down in the list
severity remove      Remove a severity value
ticket remove        Remove ticket
ticket_type add      Add a ticket type
ticket_type change   Change a ticket type
ticket_type list     Show possible ticket types
ticket_type order    Move a ticket type up or down in the list
ticket_type remove   Remove a ticket type
upgrade              Upgrade database to current version
version add          Add version
version list         Show versions
version remove       Remove version
version rename       Rename version
version time         Set version date
wiki dump            Export wiki pages to files named by title
wiki export          Export wiki page to file or stdout
wiki import          Import wiki page from file or stdin
wiki list            List wiki pages
wiki load            Import wiki pages from files
wiki remove          Remove wiki page
wiki rename          Rename wiki page
wiki replace         Replace the content of wiki pages from files (DANGEROUS!)
wiki upgrade         Upgrade default wiki pages to current version
Trac [/var/trac/devel/env]> help wiki dump
wiki dump <directory> [page] [...]
```

Export wiki pages to files named by title

Individual wiki page names can be specified after the directory. A name ending with a * means that all wiki pages starting with that prefix should be dumped. If no name is specified, all wiki pages are dumped.

```
Trac [/var/trac/devel/env]>
root@douglas:/var/trac/devel# mv ~feurig/tracpwd.old env/.htpasswd
root@douglas:/var/trac/devel# trac-admin /var/trac/devel/env/ permission add feurig TRAC_ADMIN
root@douglas:/var/trac/devel# trac-admin /var/trac/devel/env/ permission add joe TRAC_ADMIN
root@douglas:/var/trac/devel# cp ~feurig/sd_logo_sm.png env/htdocs/
root@douglas:/var/trac/devel# chmod oug+r env/htdocs/sd_logo_sm.png
root@douglas:/var/trac/devel# nano env/conf/trac.ini
```

5.60 Nigel

Nigel is a TP-Link Mr3020 router running LEDE 17.01.3 that I am using to connect things to the net. Things are connected in one of 2 ways.

- wifi
- serial
- usb-serial
- 3.3V serial.

5.60.1 Wifi

Nigel provides a Hidden wifi access point called critters on the internal lan (192.168.2.0/24)

```
root@nigel:~# nano /etc/config/wireless
config wifi-device radio0
    option type mac80211
    option channel 11
    option hwmode 11g
    option path 'platform/ar933x_wmac'
    option htmode HT20
    option disabled 0

config wifi-iface
    option device radio0
    option network lan
    option hidden 1
    option mode ap
    option ssid crl3t3rs
    option encryption psk2
    option key '*****'
root@nigel:~#
```

- [wiki:OpenWRTonMR3020 Setting up OpenWRT (15.05) on a TP-link MR3020]
- <https://downloads.lede-project.org/releases/17.01.4/targets/ar71xx/generic/>

5.61 NotesAddingAnsibleToContainerCreation

5.61.1 Adding Ansible to Container Creation

Container creation using ansible involved modifying some existing examples and tweaking things that worked using LXD and its cloud init. The result is the ability to create base containers with built in admin users and ssh key based connectivity.

The system setup and admin user installation is done by the cloud-init portions of the susdev19 lxc profile. The network configuration is passed as part of creating the container. The disk and network device was moved to a separate profile allowing containers to have different disk or network connections. Because ansible requires python the create-lxd-containers playbook waits for cloud init to finish and then checks for python and attempts to install it if its not there.

All files used for ansible as well as the susdev19 lxc profile can be checked out of the private repository.

<https://bitbucket.org/suspectdevicesadmin/ansible/src/master/>

File layout. * /etc/ansible/ * hosts/ – base inventory file * playlists/ – playlists * files/ – files (*also where we put the profiles*) * roles/

Currently this does not work to create containers on bs2020. To create a container on bs2020 * create container on kb2018 * move the container to bs2020 * adjust the host file

Things that needed to be changed in our current environment. * Profile needed to be broken out between network, system setup, and device mapping. * Network configuration is generated on the fly using a file or using ansible * User configuration and minimal software setup are now shared using the susdev19 profile * Default network devices and disk pools can be overridden using a separate profile (*infra for instance*)

- ansible vs cloud-init
- Cloud init should provide a distro agnostic way to add users, keys and software.
- Images do not always provide cloud init and even that may not be fully functional.
- Ansible allows per distro scripting but not distro agnostic modules for many tasks

5.61.2 Linkdump

- <https://blog.sourcecode.de/posts/2016/11/25/how-to-create-lxd-containers-with-ansible-2-2/>
- <https://dev.to/livioribeiro/using-lxd-and-ansible-to-simulate-infrastructure-2g8l>
- <https://medium.com/@abhijeet.kamble619/10-things-you-should-start-using-in-your-ansible-playbook-808daff76b65>
- <https://leucos.github.io/ansible-files-layout>

5.62 Notes: Automating Container Updates

This would have worked in an lxc only world...

```
#!/bin/bash
# Purpose: Update all lxc vms
# Note: Tested on Ubuntu LTS only
# Author: Vivek Gite <www.cyberciti.biz>, under GPL v2+
# -----

# Get the vm list
vms="$(lxc-ls --active)"

# Update each vm
update_vm(){
    local vm="$1"
    echo "**** [VM: $vm [$(hostname) @ $(date)] ] ****"
    /usr/bin/lxc-attach -n "$vm" apt-get -- -qq update
    /usr/bin/lxc-attach -n "$vm" apt-get -- -qq -y upgrade
    /usr/bin/lxc-attach -n "$vm" apt-get -- -qq -y clean
    /usr/bin/lxc-attach -n "$vm" apt-get -- -qq -y autoclean
    # Note for RHEL/CentOS/Fedora Linux comment above two line and uncomment the following line #
    # lxc-attach -n "$vm" yum -y update
    echo "-----"
}

# Do it
for v in $vms
do
    update_vm "$v"
done
```

This works for updating everything debian under the lxd.. Not sure you need anything else :)

```
#!/bin/bash
# A simple shell script to update all lxd container hypervisor
# URL: https://bash.cyberciti.biz/virtualization/shell-script-to-update-all-lxd-container-hypervisor/
# Tested on : Ubuntu 16.04 LTS lxd server
# Tested on : Ubuntu/Debian lxd container hypervisor only
# -----
# Author: nixCraft
# Copyright: 2016 nixCraft under GNU GPL v2.0+
# -----
# Last updated 14 Aug 2016
# -----
# Set full path to bins
_apt="/usr/bin/apt-get"
_lxc="/usr/bin/lxc"
_awk="/usr/bin/awk"

# Get containers list
clist="$({_lxc} list -c ns | ${_awk} '!/NAME/{ if ( $4 == "RUNNING" ) print $2}')"

# Use bash for loop and update all container hypervisor powered by Debian or Ubuntu
# NOTE: for CentOS use yum command instead of apt-get
for c in $clist
do
    echo "Updating Debian/Ubuntu container hypervisor \"${c}\"..."
    ${_lxc} exec $c ${_apt} -- -qq update
    ${_lxc} exec $c ${_apt} -- -qq -y upgrade
    ${_lxc} exec $c ${_apt} -- -qq -y clean
    ${_lxc} exec $c ${_apt} -- -qq -y autoclean
done
```

Shell Fragment for looking at os distribution.

```
# Determine OS platform
UNAME=$(uname | tr "[:upper:]" "[:lower:]")
# If Linux, try to determine specific distribution
if [ "$UNAME" == "linux" ]; then
    # If available, use LSB to identify distribution
    if [ -f /etc/lsb-release -o -d /etc/lsb-release.d ]; then
        export DISTR0=$(lsb_release -i | cut -d: -f2 | sed s/'^t'//)
    # Otherwise, use release info file
    else
        export DISTR0=$(ls -d /etc/[A-Za-z]*[_-][rv]e[lr]* | grep -v "lsb" | cut -d '/' -f3 | cut -d '-' -f1 | cut -d '_' -f1)
    fi
fi

# For everything else (or if above failed), just use generic identifier
[ "$DISTR0" == "" ] && export DISTR0=$UNAME
unset UNAME
```

Linkdump

- <https://askubuntu.com/questions/459402/how-to-know-if-the-running-platform-is-ubuntu-or-centos-with-help-of-a-bash-scri>
- <https://ask.fedoraproject.org/en/question/49738/how-to-check-if-system-is-rpm-or-debian-based/>
- <http://fuckingshellscripts.org/>
- <https://etbe.coker.com.au/2007/08/30/identifying-the-distribution-of-a-linux-system/>
- <https://ask.fedoraproject.org/en/question/49738/how-to-check-if-system-is-rpm-or-debian-based/>
- <https://hvops.com/articles/ansible-vs-shell-scripts/>
- <https://news.ycombinator.com/item?id=6431552>
- <https://www.cyberciti.biz/faq/how-to-update-debian-or-ubuntu-linux-containers-lxc/>
- <https://blog.sleeplessbeastie.eu/2017/08/21/how-to-upgrade-lxd-guests/>
- <https://blog.selectel.com/managing-containers-lxd-brief-introduction/>
- <http://xmodulo.com/lxc-containers-ubuntu.html>

5.63 Buster Notes

FreedomBox is packaged on Debian 10

Creating a cloud-init capable Debian/10 container

Download container from images.

```
root@annie:~# lxc image copy images:debian/10 local: --copy-aliases
root@annie:~# lxc image list
```

ALIAS	FINGERPRINT	PUBLIC	DESCRIPTION	ARCH	SIZE	UPLOAD DATE
b (10 more)	c395a7105278	no	ubuntu 18.04 LTS amd64 (release) (20180911)	x86_64	173.98MB	Sep 30, 2018 at 4:00am (UTC)
debian/10 (7 more)	ec89a28d9d81	no	Debian buster amd64 (20190930_05:24)	x86_64	73.00MB	Sep 30, 2019 at 3:58pm (UTC)

```
root@annie:~# lxc init debian/10 buster
Creating buster
root@annie:~# lxc start buster
```

Copy templates and metadata to image

```
root@annie:/var/lib/lxd/storage-pools/devil/containers/buster# cat ../viva/metadata.yaml >>metadata.yaml
root@annie:/var/lib/lxd/storage-pools/devil/containers/buster# cp -rpv ../viva/templates/
cloud-init-meta.tpl cloud-init-network.tpl cloud-init-user.tpl cloud-init-vendor.tpl hostname.tpl
root@annie:/var/lib/lxd/storage-pools/devil/containers/buster# cp -rpv ../viva/templates .
'../viva/templates/cloud-init-meta.tpl' -> './templates/cloud-init-meta.tpl'
'../viva/templates/cloud-init-network.tpl' -> './templates/cloud-init-network.tpl'
'../viva/templates/cloud-init-user.tpl' -> './templates/cloud-init-user.tpl'
'../viva/templates/hostname.tpl' -> './templates/hostname.tpl'
'../viva/templates/cloud-init-vendor.tpl' -> './templates/cloud-init-vendor.tpl'
root@annie:/var/lib/lxd/storage-pools/devil/containers/buster# nano metadata.yaml
.... delete original templates section and properties from other system ....
```

Add cloud-init, cloud-utils, and ssh server

```
root@annie:~# lxc exec buster bash
root@buster:~# apt-get install inetutils-ping nano cloud-init cloud-utils openssh-server python3
Reading package lists... Done
Building dependency tree
Reading state information... Done
...
root@buster:~# nano /etc/network/interfaces.d/50-cloud-init.cfg
root@buster:~# rm /etc/network/interfaces
root@buster:~# ln -s /etc/network/interfaces.d/50-cloud-init.cfg /etc/network/interfaces

root@buster:~# shutdown -h now
```

Publish the image

```
root@annie:~# lxc publish buster --alias debian/10cloud description="Debian buster plus cloud-init"
```

5.63.1 Using the container

If you are not going to keep the image you can create it using lxc init.

```
root@annie:/etc/ansible# lxc init debian/10cloud camo -p susdev19 -p default
root@annie:/etc/ansible# lxc start camo
```

Or you can add it to /etc/ansible/hosts and use the create-lxc-containers.yml playbook.

```
root@annie:/etc/ansible# grep camo hosts
camo ip_address=192.168.0.253 purpose="Freedombox Test Server" image_alias=debian/10cloud
root@annie:/etc/ansible# ansible-playbook playbooks/create-lxd-containers.yml
...
root@annie:/etc/ansible# lxc list
```

NAME	STATE	IPV4	IPV6	TYPE	SNAPSHOTS
buster	STOPPED			PERSISTENT	0
camo	RUNNING	192.168.0.253 (eth0)		PERSISTENT	0

5.63.2 Installing freedombox

```
# apt-get install freedombox
```

There are a few questions on the install that need to be answered and then its more or less done. I am not sure I want it exposed until I figure out how to configure it securely. Am going to run it up on the home server first.

5.64 HP Z400 notes

What to do when you get a desktop version of ubuntu and you want a server

```

sudo bash
apt-get update&&apt-get dist-upgrade && apt-get autoremove
nano /etc/hostname

root@joey:~# cd /etc/netplan/
root@joey:~# mv 01-network-manager-all.yaml /tmp/
root@joey:~# nano 50-cloud-init.yaml
network:
  version: 2
  renderer: networkd
  ethernets:
    enp1s0:
      dhcp4: no
      dhcp6: no
  bridges:
    br0:
      dhcp4: no
      dhcp6: no
      addresses:
        - 192.168.0.65/24
      gateway4: 192.168.0.1
      nameservers:
        addresses:
          - 192.168.0.1
          - 198.202.31.141
  interfaces:
    - enp1s0

root@joey:~# netplan apply
root@joey:~# reboot

root@joey:~# apt-get install openssh-server

root@joey:~# nano /etc/default/grub
root@joey:~# update-grub
root@joey:~# systemctl enable multi-user.target --force
root@joey:~# systemctl set-default multi-user.target
root@joey:~# reboot

```

5.64.1 Memory

Looking at the memory usage for that standalone desktop it seems like we could get by with 4G until we decide to use zfs or containers. Our file server is running 3 containers and 7T of zfs and uses 12.8G.(out of 21)

The cheapest memory I could find at present is on eBay at \$4/G [https://www.ebay.com/itm/NEW-8GB-2x4GB-Memory-RAM-PC3-10600-ECC-Unbuffered-HP-Compaq-Workstation-Z400/221132609031?](https://www.ebay.com/itm/NEW-8GB-2x4GB-Memory-RAM-PC3-10600-ECC-Unbuffered-HP-Compaq-Workstation-Z400/221132609031?ssPageName=STRK%3AMEBIDX%3AIT&_trksid=p2057872.m2749.l2649)

On Board "Fake" Raid

The bios raid on the hpz400 is a bastardized spitwad created by intel hp and only seems to play well with micro\$oft. The drivers which kind of worked under dmraid have been integrated into linux's mdraid drivers. The newfangled installer on 18.04 server (subiquity) breaks on it. The alternative installer installs an os on the raid array but it won't boot. You can boot it from another disk but it won't boot by itself.

MAKING IT WORK

Since I purchased a 240G ssd for this server I installed ubuntu manually to a 10G partition on it, When installing the os activate the intel raid but not the sata raid. Rebooting the box gives you the option of booting to the bios raid array. From there you can chroot onto the boot disk and update the grub default to boot the other disk.

```

root@joey:~# mount /dev/sdd1 /mnt
root@joey:~# mount -t proc proc /mnt/proc
root@joey:~# mount -t sysfs sys /mnt/sys
root@joey:~# mount -o bind /dev /mnt/dev
root@joey:~# chroot /mnt
root@joey:/# cd etc

```

```
root@joey:/etc# nano default/grub
...
GRUB_DEFAULT=2
root@joey:/etc# update-grub
```

On the other hand

https://www.newegg.com/Product/Product.aspx?Item=9SIAC0F8UV0008&ignorebbr=1&source=region&nm_mc=KNC-GoogleMKP-PC&cm_mmc=KNC-GoogleMKP-PC--pla-PC+Server+and+Parts-Hard+Drive+Controllers+%2F+RAID+Cards-_9SIAC0F8UV0008&gclid=CjwKCAjwza_mBRBTEiwASDWVvrfrvrffm4o0noMgtv3UEV7bdZpf1JgLYx4v99kFnn_iNhMc7H2bPhoC7YoQAvl
 " Note: There may be a bios upgrade to fix this however z400 will not boot from SmartArray P810 either.

Processor

System !#1

```
root@annie:~# dmesg|grep smp
[ 0.000000] smpboot: Allowing 16 CPUs, 8 hotplug CPUs
[ 0.044000] smpboot: CPU0: Intel(R) Xeon(R) CPU W3520 @ 2.67GHz (family: 0x6, model: 0x1a, stepping: 0x5)
[ 0.044000] smp: Bringing up secondary CPUs ...
[ 0.062447] smp: Brought up 1 node, 8 CPUs
[ 0.062447] smpboot: Max logical packages: 2
[ 0.064004] smpboot: Total of 8 processors activated (42670.64 BogoMIPS)
```

System !#2

```
root@joey:~# dmesg|grep smp
[ 0.000000] smpboot: Allowing 16 CPUs, 14 hotplug CPUs
[ 0.044000] smpboot: CPU0: Intel(R) Xeon(R) CPU W3503 @ 2.40GHz (family: 0x6, model: 0x1a, stepping: 0x5)
[ 0.044000] smp: Bringing up secondary CPUs ...
[ 0.046748] smp: Brought up 1 node, 2 CPUs
[ 0.046748] smpboot: Max logical packages: 8
[ 0.046748] smpboot: Total of 2 processors activated (9599.50 BogoMIPS)
root@joey:~#
```

Same after upgrade to 1xX5650 \$10 on eBay.

```
root@DeeDee:~# dmesg|grep smp
[ 0.000000] smpboot: Allowing 16 CPUs, 10 hotplug CPUs
[ 0.044000] smpboot: CPU0: Intel(R) Xeon(R) CPU X5650 @ 2.67GHz (family: 0x6, model: 0x2c, stepping: 0x2)
[ 0.044000] smp: Bringing up secondary CPUs ...
[ 0.060027] smp: Brought up 1 node, 6 CPUs
[ 0.060027] smpboot: Max logical packages: 3
[ 0.060027] smpboot: Total of 6 processors activated (32002.70 BogoMIPS)
```

Compared to BS2020:

```
root@bs2020:~# dmesg|grep smp
[ 0.000000] smpboot: Allowing 24 CPUs, 0 hotplug CPUs
[ 0.164444] smpboot: CPU 0 Converting physical 1 to logical package 0
[ 0.172000] smpboot: CPU0: Intel(R) Xeon(R) CPU X5650 @ 2.67GHz (family: 0x6, model: 0x2c, stepping: 0x2)
[ 0.216015] smp: Bringing up secondary CPUs ...
[ 0.004000] smpboot: CPU 1 Converting physical 0 to logical package 1
[ 0.474623] smp: Brought up 2 nodes, 24 CPUs
[ 0.480002] smpboot: Max logical packages: 2
[ 0.484002] smpboot: Total of 24 processors activated (127681.26 BogoMIPS)
```

and kb2018

```
root@kb2018:~# dmesg|grep smp
[ 0.000000] smpboot: Allowing 32 CPUs, 16 hotplug CPUs
[ 0.140000] smpboot: CPU0: Intel(R) Xeon(R) CPU E5640 @ 2.67GHz (family: 0x6, model: 0x2c, stepping: 0x2)
[ 0.176015] smp: Bringing up secondary CPUs ...
[ 0.342596] smp: Brought up 2 nodes, 16 CPUs
[ 0.346074] smpboot: Max logical packages: 4
[ 0.348004] smpboot: Total of 16 processors activated (85310.05 BogoMIPS)
```

Upgrading the Processor

According to [this thread](#) the z400 with the 6 memory slots will support the Xeon 56xx family. <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/xeon-5600-brief.pdf>

I purchased one on eBay for 10 bucks. I probably should have purchased 2

<https://www.ebay.com/itm/163034026449>

5.65 NotesOnAppleTalk3vsUbuntu

5.66 DL380 Raid Bios notes

5.66.1 Configuring the disks using the raid controller bios

... use some words here ... "Also reattach the images"

```
steve:~ don$ ssh -p 22222 feurig@vpn.suspectdevices.com
User:feurig logged-in to kb2018.suspectdevices.com(192.168.31.119 / FE80::9E8E:99FF:FE0C:BAD8)
iLO 3 Advanced for BladeSystem 1.88 at Jul 13 2016
Server Name: kb2018
Server Power: On

hpiLO-> vsp


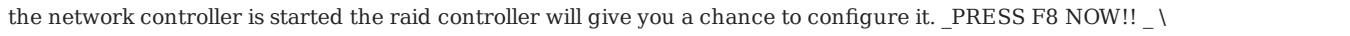

Virtual Serial Port Active: COM2

Starting virtual serial port.
Press 'ESC (' to return to the CLI Session.

root@kb2018:~# fdisk -l|grep Disk\ \
Disk /dev/loop0: 86.9 MiB, 91099136 bytes, 177928 sectors
Disk /dev/loop1: 87.9 MiB, 92164096 bytes, 180008 sectors
Disk /dev/loop2: 63.4 MiB, 66486272 bytes, 129856 sectors
Disk /dev/sda: 136.7 GiB, 146778685440 bytes, 286677120 sectors

root@kb2018:~# reboot
[ OK ] Stopped Stop ureadahead data collection 45s after completed      Stopping Session 98 of user feurig.
      Stopping Availability of block devices...
...

[ OK ] Reached target Shutdown.
[ OK ] Reached target Final Step.
      Starting Reboot...
[292357.910620] reboot: Restarting system
```

After several seconds you will see a text based bios screen  After the network controller is started the raid controller will give you a chance to configure it.  After the network controller is started the raid controller will give you a chance to configure it. 

If you miss it you will have to escape back to the ILO3 and power cycle the machine. *(This is ok because the disks are not active until the machine actually boots)*

```
Booting from Hard Drive C:
<ESC> (
hpiLO-> power off hard

status=0
status_tag=COMMAND COMPLETED
Wed Sep 26 15:31:57 2018

Forcing server power off .....
Please wait 6 seconds for this operation to complete.

hpiLO-> power

status=0
status_tag=COMMAND COMPLETED
Wed Sep 26 15:32:04 2018

power: server power is currently: Off

hpiLO-> power on

status=0
status_tag=COMMAND COMPLETED
Wed Sep 26 15:32:21 2018

Server powering on .....

hpiLO-> vsp

Virtual Serial Port Active: COM2
```

```
Starting virtual serial port.
Press 'ESC (' to return to the CLI Session.
```

Once in the raid controller bios you will get a main menu.

[[Image(CaptiveRaidController:ViewLogicalDrive.png)]]

If you select view logical drives will see that the first two disks are combined into a mirrored pair and that there are no other drives defined.

So we select "Create Logical Drive". Which gives us the following screen.

[[Image(CaptiveRaidController:CreateLogicalDriveDefaults.png)]]

Notice that the defaults are to create a raid 1+0 array with the first two matching disks. Deselecting either disk (down arrow, spacebar) will cause the raid configuration to automatically drop to RAID 0

Press Enter when finished. The next screen will ask you to verify the creation

Repeat this for each remaining disk.

When you are finished you can view the logical drives. [[Image(CaptiveRaidController:RaidConfFinished.png)]]

The key will walk you back out so you can continue to boot.

5.66.2 success

```
root@kb2018:~# fdisk -l|grep Disk\ \
Disk /dev/loop0: 86.9 MiB, 91099136 bytes, 177928 sectors
Disk /dev/loop1: 87.9 MiB, 92164096 bytes, 180008 sectors
Disk /dev/loop2: 63.4 MiB, 66486272 bytes, 129856 sectors
Disk /dev/sda: 136.7 GiB, 146778685440 bytes, 286677120 sectors
Disk /dev/sdb: 223.6 GiB, 240021504000 bytes, 468792000 sectors
Disk /dev/sdc: 223.6 GiB, 240021504000 bytes, 468792000 sectors
Disk /dev/sdd: 279.4 GiB, 299966445568 bytes, 585871964 sectors
Disk /dev/sde: 279.4 GiB, 299966445568 bytes, 585871964 sectors
root@kb2018:~#
```


5.67 ILO3 Notes

The ILO 3 card on the HP ProLiant DL380 allows us complete remote control of the server for this reason the same security precautions which are used on the idrac6 need to be implemented.

Securing the ILO3

The ilo3 is not directly accessible except through the admin lan firewall. Eventually this will require vpn access however in the mean time it is accessed through port redirection. The ilo3s main access is through https. The port number for this is configurable along with the other ports used. (ssh + 2 ports for console redirection)

[[Image(ILO3Notes:ilo3NetworkPorts.png)]] Unless you are working in a MAAS environment the ipv6 should be disabled and the ipv4 address should be made static. This will require resetting the ILO3 itself. [[Image(ILO3Notes:ILO3ResetILO.png)]]

MANAGE ADMIN ACCOUNTS

Create user and management accounts as soon as possible and demote or remove any existing accounts.

[[Image(ILO3Notes:ilo3UserAdmin.png)]] While there you should add your ssh keys for ssh connections. Note that only dsa keys are supported so you may need to create a separate public key.

```
steve:~ don$ ssh-keygen -t dsa
Generating public/private dsa key pair.
...
```

Java Console

The ILO 3 provides a java console similar to the one provided by the Dell idrac. It requires the remote console port (17990) as well as the Virtual Media Port (17988) to function properly. [[Image(ILO3Notes:HPBootSplash.png)]]

Remote Media

Attaching an iso is straight forward. [[Image(ILO3Notes:ilo3RemovableMedia.png)]] Using the Ubuntu 18.04 Live Server over a DSL connection is pokey and complains a lot but it does not fail. [[Image(ILO3Notes:ilo3NetworkMountsAndLag.png)]]

Enabling bios and console access via ssh.

Once you have administrative access to the ILO3 and you have an os install you can do everything vial ssh. Much like the idrac you need access to the f9 key. [[Image(wiki:Idrac6:fnkeys.png)]] * Enter bios * Select Serial settings. * set console redirection to com2 _ you will have to do this in the advanced settings as well _ [[Image(ILO3Notes:ILO Bios Virtual Serial Port.jpg)]]

5.67.1 Connecting to the console

Once the bios is set up you can ssh to the console using your iso credentials and ssh key.

```
steve:~ don$ ssh -p22222 feurig@vpn.suspectdevices.com
User:feurig logged-in to kb2018.suspectdevices.com(192.168.31.119 / FE80::9E8E:99FF:FE0C:BAD8)
iLO 3 Advanced for BladeSystem 1.88 at Jul 13 2016
Server Name: kb2018
Server Power: On

hpiLO-> help

status=0
status_tag=COMMAND COMPLETED
Sat Sep 22 20:20:42 2018
...
DMTF SMASH CLP Commands:
...
HP CLI Commands:

POWER      : Control server power.
UID        : Control Unit-ID light.
NMI        : Generate an NMI.
VM         : Virtual media commands.
LANGUAGE   : Command to set or get default language
```

```
VSP      : Invoke virtual serial port.
TEXTCONS : Invoke Remote Text Console.
```

Then you can connect to the console

```
hpiLO-> vsp

Virtual Serial Port Active: COM2

Starting virtual serial port.
Press 'ESC (' to return to the CLI Session.

Ubuntu 18.04.1 LTS kb2018 ttyS1

kb2018 login:
```

If the session is preoccupied use the following (stop /system1/oemhp_vsp1)

```
steve:~ don$ ssh -p 22222 feurig@vpn.suspectdevices.com
User:feurig logged-in to kb2018.suspectdevices.com(192.168.31.119 / FE80::9E8E:99FF:FE0C:BAD8)
iLO 3 Advanced for BladeSystem 1.88 at Jul 13 2016
Server Name: kb2018
Server Power: On

hpiLO-> vsp
...
Virtual Serial Port is currently in use by another session.
hpiLO-> stop /system1/oemhp_vsp1
...
hpiLO-> hpiLO-> vsp

Virtual Serial Port Active: COM2
```



5.67.2 fixing grub (identical to the process for idrac 6)

You need set the console to ttyS1 by adding a console=ttyS1,115200n8 to the end of the kernel line

```
root@bs2020:~# nano /boot/grub/menu.list
...
kernel                /boot/vmlinuz-4.4.0-96-generic root=UUID=8cafbd6f-441e-4f76-b89c-017fc22253f9 ro console=hvc0 console=ttyS1,115200n8
```

Add the changes to /etc/default/grub so that it will survive updates to the kernel.

```
root@bs2020:~# nano /etc/default/grub
...
GRUB_TERMINAL='serial console'
GRUB_CMDLINE_LINUX="console=hvc0 console=ttyS1,115200n8"
GRUB_SERIAL_COMMAND='serial --speed=115200 --unit=1 --word=8 --parity=no --stop=1'
root@bs2020:~# update-grub
```

Reboot the server and attach to the console.  

5.67.3 virtual serial port in action

In order to make the dl380 expose the disks we added required jumping into the raid controllers bios during boot and configuring it. This is documented [\[\[wiki:CaptiveRaidController|here\]\]](#)

HP Documents

- [ILO3 Users Guide \(https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-c02774507\)](https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-c02774507)
- [ILO3 Scripting Guide \(https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-c02774508\)](https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-c02774508)
- [ILO3 Serial Port Guide \(https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-c00263709\)](https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-c00263709)
- [ILO3 Security Brief \(https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-a00026171en_us\)](https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=5294355&docLocale=en_US&docId=emr_na-a00026171en_us)

Link Dump


- [using the VSP features of the ilo3 to configure the raid controller \(http://trac.suspectdevices.com/trac/wiki/CaptiveRaidController\)](http://trac.suspectdevices.com/trac/wiki/CaptiveRaidController)
- [Using IPMI to configure ILO card \(http://dev-random.net/configuring-hp-ilo-through-linux-automatically/\)](http://dev-random.net/configuring-hp-ilo-through-linux-automatically/)
- <https://sysadmin.compextreme.ro/access-hps-ilo-remote-console-via-ssh/>
- [bonus link on how to kill outstanding connections \(https://stivesso.blogspot.com/2012/02/hp-ilolinux-output-to-vsp-for-linux.html\)](https://stivesso.blogspot.com/2012/02/hp-ilolinux-output-to-vsp-for-linux.html)

5.68 Irac6 Notes

The Dell idrac is a very powerful tool allowing remote administration of a server down to bare bones os installation. The console feature of this tool is based on a Java app which is downloaded from the idrac and which then sets up a vnc style remote console. As the hardware ages this code becomes less and less secure and is often broken but updates to the local os (OS X being ours) and to java.

When purchasing a dell with idrac capabilities ALWAYS opt for the "Enterprise" edition. * The enterprise edition uses a separate network connection allowing it to be placed on a secure lan. * The enterprise edition allows remote console

5.68.1 securing the idrac

- Since the idrac allows complete control of the system it should never be allowed directly on the network.
- The idrac is initially configured with a "root" user who's password is "calvin"
- Change it ASAP
- Create local administrators.
- Once local accounts are tested strip the root user of all privileges. 
- The idrac needs several ports opened to be controlled.
- https
- ssh
- vnc (5900)


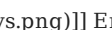


5.68.2 remote console fun

The latest version of java disables the graphical remote console. The console makes you jump through all of the hoops to run launches and fails to connect. The solution is to disable the disabling of ssl3.

```
bash-3.2# find / -name java.security -print
/Applications/Arduino.app/Contents/PlugIns/JavaAppletPlugin.plugin/Contents/Home/lib/security/java.security
/Applications/microchip/mplabx/v3.15/sys/java/jre1.7.0_79.jre/Contents/Home/lib/security/java.security
/Applications/microchip/mplabx/v3.15/sys/java/jre1.8.0_60.jre/Contents/Home/lib/security/java.security
/Applications/Xcode.app/Contents/Applications/Application Loader.app/Contents/itms/java/lib/security/java.security
find: /dev/fd/3: Not a directory
find: /dev/fd/4: Not a directory
/Library/Internet Plug-Ins/JavaAppletPlugin.plugin/Contents/Home/lib/security/java.security
/Users/don/Downloads/Energia.app/Contents/PlugIns/jdk1.8.0_91.jdk/Contents/Home/jre/lib/security/java.security
bash-3.2# nano /Library/Internet Plug-Ins/JavaAppletPlugin.plugin/Contents/Home/lib/security/java.security
....
... change the commented out line to the one below it ...
#jdk.tls.disabledAlgorithms=SSLv3, RC4, MD5withRSA, DH keySize < 1024, \
jdk.tls.disabledAlgorithms, RC4, MD5withRSA, DH keySize < 1024, \
    EC keySize < 224, DES40_CBC, RC4_40, 3DES_EDE_CBC
....
```

5.68.3 Enabling bios and console access via ssh.

...Unfortunately this requires a functioning console.. As well as access to the f2 key.

 Enter bios  Select Serial settings.  set console redirection to com2 

5.68.4 Connecting to the console

Once you can ssh to the idrac set up the serial using racadm

```
steve:~ don$ ssh -p222 feurig@198.202.31.242
feurig@198.202.31.242's password:
/admin1-> racadm config -g cfgSerial -o cfgSerialBaudRate 115200
Object value modified successfully
/admin1-> racadm config -g cfgSerial -o cfgSerialCom2RedirEnable 1
Object value modified successfully
```

```
/admin1-> racadm config -g cfgSerial -o cfgSerialSshEnable 1
Object value modified successfully
/admin1-> racadm config -g cfgIpmiSol -o cfgIpmiSolEnable 1
Object value modified successfully
/admin1-> racadm config -g cfgIpmiSol -o cfgIpmiSolBaudRate 115200
Object value modified successfully
```

Then you can connect to the console

```
/admin1-> console com2
```

```
Connected to Serial Device 2. To end type: ^\
```

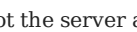
5.68.5 fixing grub

You need set the console to ttyS1 by adding a console=ttyS1,115200n8 to the end of the kernel line

```
root@bs2020:~# nano /boot/grub/menu.list
...
kernel          /boot/vmlinuz-4.4.0-96-generic root=UUID=8cafbd6f6-441e-4f76-b89c-017fc22253f9 ro console=hvc0 console=ttyS1,115200n8
```

Add the changes to /etc/default/grub so that it will survive updates to the kernel.

```
root@bs2020:~# nano /etc/default/grub
...
GRUB_TERMINAL='serial console'
GRUB_CMDLINE_LINUX="console=hvc0 console=ttyS1,115200n8"
GRUB_SERIAL_COMMAND="serial --speed=115200 --unit=1 --word=8 --parity=no --stop=1"
root@bs2020:~# update-grub
```

Reboot the server and attach to the console. 



5.68.6 as of Ubuntu 16.04 systemd actually figures it out from there

Once the console is set systemd creates a getty process for it. Otherwise you can chase the web around for the pre-upstart (i.e. /etc/inittab), upstart (/etc/init/xxxx), and early systemd questions/solutions. Hope they don't screw it up in 18.04 * Victory!!!



5.68.7 Afterthoughts / Todo

- look at ipmi <http://www.alleft.com/sysadmin/ipmi-sol-inexpensive-remote-console/>

5.68.8 linkdump

- http://media.community.dell.com/en/dtc/attach/idrac6_security_v1.pdf
- <https://gist.github.com/xbb/4fd651c2493ad9284dbcb827dc8886d6>
- <https://www.dell.com/community/Systems-Management-General/iDRAC6-Virtual-Console-Connection-Failed/td-p/5144021/page/3>
- <http://support.hkti.net/support/solutions/articles/3000003121-for-dell-user-how-to-open-dell-idrac-virtual-console>
- <https://www.slac.stanford.edu/grp/cd/soft/unix/EnableSerialConsoleAccessViaSSH.htm>
- <https://serverfault.com/questions/269382/garbled-using-from-dell-drac-for-serial-console-redirection>
- <https://www.serverhome.nl/media/specsheets/Dell/DRAC/iDRAC6-user-guide.pdf>
- <http://jonamiki.com/2014/10/18/sol-serial-over-lan-connection-from-linux-to-dell-idrac-or-bmc/>
- <https://www.hiroom2.com/2016/06/06/ubuntu-16-04-grub2-and-linux-with-serial-console/>
- <http://0pointer.de/blog/projects/serial-console.html>
- <https://lnxgeek.wordpress.com/2018/02/16/serial-console-howto-ubuntu-16-04/>
- <http://lukeluo.blogspot.com/2015/04/dell-r710-idrac6-setup-with-ssh-console.html>

5.69 LXD FIRST IMPRESSIONS:

Creating LXD Container with Static IP (and Docker Profile) We want to create a docker capable LXD container using an existing bridge with a static ip and zfs. Then we want to install docker and test it. We will make a copy of this container once the admin users have been added so that we wont have to replicate these tasks. Our security model requires ssh keys to log in AND passwords to escalate privileges.

The first thing we learned is that LXC and LXD are pretty different beasts and that while lxc with lxc-templates is a straightforward way to create containers that act a lot like regular old hardware LXD brings on all of its we love the mother fucking cloud baggage. Major differences had to do with user mapping on the containers files created by root on the host were mapped to nobody on the container, making it really difficult to set up home directories etc. (for a workaround to this see <https://stackoverflow.com/questions/33377916/migrating-lxc-to-lxd>) The second was the way that the network is initialized with the assumption that LXD would be providing the bridge and the context.

5.69.1 First Attempt and zfs/bridge setup

Create zfs container and bridge as before

```
root@bs2020:~# lxd init
... create new zfs pool and use all of /dev/sdd1 do not configure bridge ...
root@bs2020:~# dpkg-reconfigure -p medium lxd
... no yes brl ... use existing bridge...
root@bs2020:~# lxc launch ubuntu:16.04 franklin -p default -p docker
root@bs2020:~# lxc stop franklin
root@bs2020:~# passwd -l ubuntu -R /var/lib/lxd/containers/franklin.zfs/rootfs
passwd: user 'ubuntu' does not exist
root@bs2020:~# cd /var/lib/lxd/containers/franklin.zfs/rootfs/
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat ~feurig/passwd.add>>etc/passwd
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat ~feurig/shadow.add>>etc/shadow
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# tar -xzf ~feurig/fnj.tgz
home/feurig
...
home/joe/.ssh/authorized_keys
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# usermod -R /var/lib/lxd/containers/franklin.zfs/rootfs -G sudo,root joe
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# usermod -R /var/lib/lxd/containers/franklin.zfs/rootfs -G sudo,root feurig
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# groupadd -R /var/lib/lxd/containers/franklin.zfs/rootfs -g 1001 feurig
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# groupadd -R /var/lib/lxd/containers/franklin.zfs/rootfs -g 1002 joe
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat <<eod >>/var/lib/lxd/containers/franklin.zfs/rootfs/etc/resolvconf/resolv.conf.d/base
> dns-nameserver 198.202.31.132 198.202.31.141 8.8.8.8
> nameserver 198.202.31.132 198.202.31.141 8.8.8.8
> eod
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# sed -i 's/^iface eth0/#iface eth0/' /var/lib/lxd/containers/franklin.zfs/rootfs/etc/network/interfaces
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# cat <<eod2 >>/var/lib/lxd/containers/franklin.zfs/rootfs/etc/network/interfaces
> iface eth0 inet static
>     address 198.202.31.201/25
>     gateway 198.202.31.129
>     dns-nameservers 198.202.31.132 198.202.31.141 8.8.8.8
>     dns-search suspectdevices.com digithink.com
> eod2
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# lxc start franklin
```

Try to log in to instance over the network..... FAIL Unlike lxc's ubuntu:16.04, lxd's ubuntu:16.04 has all of the cloud cruft . That and all of the modifications to the containers directory was rootsquashed (rendering it useless).

```
root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# lxc exec franklin bash
root@franklin:~# nano /etc/network/interfaces
root@franklin:~# cat /etc/network/interfaces
# This file describes the network interfaces available on your system
# and how to activate them. For more information, see interfaces(5).

# The loopback network interface
auto lo
iface lo inet loopback

# Source interfaces
# Please check /etc/network/interfaces.d before changing this file
# as interfaces may have been defined in /etc/network/interfaces.d
# See LP: #1262951
source /etc/network/interfaces.d/*.cfg
iface eth0 inet static
    address 198.202.31.201/25
    gateway 198.202.31.129
    dns-nameservers 198.202.31.132 198.202.31.141 8.8.8.8
    dns-search suspectdevices.com digithink.com
```

First thought: remove all of the cloud crap...

```
root@franklin:~# apt-get remove cloud*
...
The following packages will be REMOVED:
  cloud-guest-utils cloud-init cloud-initramfs-copymods cloud-initramfs-dyn-netconf
0 upgraded, 0 newly installed, 4 to remove and 0 not upgraded.
After this operation, 1682 kB disk space will be freed.
Do you want to continue? [Y/n] Y
...
root@franklin:~# nano /etc/network/interfaces
... add auto eth0
root@franklin:~# reboot
root@franklin:~# root@bs2020:/var/lib/lxd/containers/franklin.zfs/rootfs# lxc list
+-----+-----+-----+-----+-----+-----+
| NAME   | STATE | IPV4   | IPV6   | TYPE   | SNAPSHOTS |
+-----+-----+-----+-----+-----+-----+
| franklin | RUNNING | 198.202.31.201 (eth0) | | PERSISTENT | 0 |
+-----+-----+-----+-----+-----+-----+
```

Second thought: Fuck that! Make it work!

5.69.2 Second attempt

(create LXD profile for suspect devices development).

```
root@bs2020:~# lxc stop franklin
root@bs2020:~# lxc delete franklin
root@bs2020:~# lxc profile create susdev
root@bs2020:~# lxc profile edit susdev
...
```

Repeat until you have a working system that can be logged into remotely

Create docker container container

```
root@bs2020:~# lxc profile show susdev
config:
  user.network_mode: link-local
  user.user-data: |
    #cloud-config
    timezone: America/Vancouver
  users:
    - name: feurig
      passwd: "... SUBSTITUTE REAL PASSWORD HASH HERE ..."
      gecos: Donald Delmar Davis
      ssh-authorized-keys:
        - ssh-rsa ... SUBSTITUTE REAL KEY HERE ... don@viscious
      groups: sudo,root
      shell: /bin/bash
    - name: joe
      passwd: "... SUBSTITUTE REALPASSWORD HASH HERE ..."
      gecos: Joseph Wayne Dumoulin
      ssh-authorized-keys:
        - ssh-rsa ...SUBSTITUTE REAL KEY HERE... jdumoulin@joeslaptop
      groups: sudo,root
      shell: /bin/bash
  manage_resolv_conf: true
  resolv_conf:
    nameservers: ['198.202.31.141', '198.202.31.132', '8.8.8.8']
    searchdomains:
      - suspectdevices.com
      - digithink.com
    domain: suspectdevices.com
    options:
      rotate: true
      timeout: 1
  write_files:
    # Set static IP address could not get this to work the "right" way
    - path: /etc/network/interfaces
      permissions: '0644'
      owner: root:root
      content: |
        auto lo
        iface lo inet loopback
        auto eth0
        # change this after first instantiation
        iface eth0 inet static
          address 198.202.31.200
          broadcast 198.202.31.255
          netmask 255.255.255.128
          gateway 198.202.31.129
          dns-nameservers 198.202.31.141 198.202.31.132 8.8.8.8
  runcmd:
    # sudo needs to be able to resolve itself to authenticate users
    # and the users are locked by default
```

```

- sed -i "s/^127.0.0.1/#127.0.0.1/" /etc/hosts
- echo 127.0.0.1 `hostname` localhost >>/etc/hosts
- passwd joe -u
- passwd feurig -u
description: Try to create a sane environment for cloud-init based operating systems
devices:
  eth0:
    name: eth0
    nictype: bridged
    parent: br1
    type: nic
name: susdev
root@bs2020:~#

root@bs2020:~# lxc list
+-----+-----+-----+-----+-----+-----+
| NAME | STATE | IPV4 | IPV6 | TYPE | SNAPSHOTS |
+-----+-----+-----+-----+-----+-----+
| test13 | RUNNING | 198.202.31.200 (eth0) | | PERSISTENT | 0 |
+-----+-----+-----+-----+-----+-----+
root@bs2020:~# lxc init ubuntu16 franklin -p susdev -p docker
Creating franklin
root@bs2020:~# lxc start franklin
root@bs2020:~# lxc exec franklin bash
root@franklin:~# tail -2 /etc/shadow
feurig:<HASHED PASSWORD>:17453:0:99999:7:::
joe:<HASHED PASSWORD>:17453:0:99999:7:::
root@franklin:~# nano /etc/network/interfaces
root@franklin:~# cat /etc/network/interfaces
auto lo
iface lo inet loopback
auto eth0
# change this after first instantiation
iface eth0 inet static
  address 198.202.31.201
  broadcast 198.202.31.255
  netmask 255.255.255.128
  gateway 198.202.31.129
  dns-nameservers 198.202.31.141 198.202.31.132 8.8.8.8
root@franklin:~# cat /etc/hosts
#127.0.0.1 localhost

# The following lines are desirable for IPv6 capable hosts
::1 ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts
127.0.0.1 franklin localhost
root@franklin:~# reboot
root@bs2020:~# lxc list
+-----+-----+-----+-----+-----+-----+
| NAME | STATE | IPV4 | IPV6 | TYPE | SNAPSHOTS |
+-----+-----+-----+-----+-----+-----+
| franklin | RUNNING | 198.202.31.201 (eth0) | | PERSISTENT | 0 |
+-----+-----+-----+-----+-----+-----+
| test13 | RUNNING | 198.202.31.200 (eth0) | | PERSISTENT | 0 |
+-----+-----+-----+-----+-----+-----+
root@bs2020:~#

```

5.69.3 References

- <http://www.whiteboardcoder.com/2016/04/cloud-init-nocloud-with-url-for-meta.html>
- <https://stgraber.org/2016/03/11/xd-2-0-blog-post-series-012/>
- <https://github.com/lxc/lxd/blob/master/doc/cloud-init.md>
- <http://www.mattjarvis.org.uk/post/lxd-openstack-cloudinit-pt1/>
- <https://sdgsystems.com/blog/understanding-and-using-lxc-and-lxd>
- <http://cloudinit.readthedocs.io/en/latest/topics/examples.html>
- <http://cloudinit.readthedocs.io/en/latest/topics/debugging.html>

5.70 NotesOnUbuntu18.04

5.70.1 Netplan / Networkd

Given the success of systemd the kids decided that they needed to rewrite the networking core using a yaml file under `/etc/netplan/` and various "renderers". If it all gets too much you can replace it with the legacy system `ifupdown` and continue to edit `/etc/network/interfaces`, etc.

```
apt-get install ifupdown
```

Otherwise read the notes to follow.

See: [Netplan Documentation \(https://netplan.io/\)](https://netplan.io/)

Static Networking with Netplan

Assuming that your cloud configuration does not overwrite it the following file produces a static ip.

```
oot@phillip:~# cat /etc/netplan/50-cloud-init.yaml
network:
  version: 2
  ethernets:
    eth0:
      dhcp4: no
      addresses: [198.202.31.223/25]
      gateway4: 198.202.31.129
      nameservers:
        search: [suspectdevices.com fromhell.com vpn]
        addresses: [198.202.31.141]
```

Bridge Networking with Netplan

```
root@annie:~# nano /etc/netplan/01-netcfg.yaml
network:
  version: 2
  renderer: networkd
  ethernets:
    ens6:
      dhcp4: true
      dhcp6: no
    enpls0:
      dhcp4: no
      dhcp6: no
  bridges:
    br0:
      dhcp4: no
      dhcp6: no
      addresses:
        - 192.168.0.66/24
      gateway4: 192.168.0.1
      nameservers:
        addresses:
          - 192.168.0.1
          - 198.202.31.141
      interfaces:
        - enpls0
root@annie:~# netplan apply
root@annie:~# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
...
2: enpls0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master br0 state UP group default qlen 1000
    link/ether 78:e7:d1:c3:ef:9e brd ff:ff:ff:ff:ff:ff
3: ens6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether 00:14:d1:25:2b:bc brd ff:ff:ff:ff:ff:ff
    inet 192.168.2.66/24 brd 192.168.2.255 scope global dynamic ens6
        valid_lft 43163sec preferred_lft 43163sec
    inet6 fd5b:alad:aeeb::fd0/128 scope global noprefixroute
...
6: br0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether aa:18:c9:5a:76:d6 brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.66/24 brd 192.168.0.255 scope global br0
        valid_lft forever preferred_lft forever
    inet6 fe80::a818:c9ff:fe5a:76d6/64 scope link
        valid_lft forever preferred_lft forever
root@annie:~# brctl show
```

```

bridge name bridge id      STP enabled interfaces
br0      8000.aa18c95a76d6  no      enpls0
root@annie:~#

```

5.70.2 And it works for anonymous bridges EXCEPT FOR THE BUG

Basically if no address is given for a bridge netplan fails to tell systemd to up the interface anyway and the bridges do not come up.

```

root@bs2020:~# nano /etc/netplan/50-cloud-init.yaml
# This file is generated from information provided by
# the datasource.  Changes to it will not persist across an instance.
# To disable cloud-init's network configuration capabilities, write a file
# /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg with the following:
# network: {config: disabled}
network:
  version: 2
  renderer: networkd
  ethernets:
    eno1:
      dhcp4: no
      addresses: [192.168.31.158/24]
      gateway4: 192.168.31.1
      nameservers:
        search: [suspectdevices.com fromhell.com vpn]
        addresses: [198.202.31.141]
    eno2:
      dhcp4: no
      optional: true
    eno3:
      dhcp4: no
    eno4:
      dhcp4: no
  bridges:
    br0:
      dhcp4: no
      dhcp6: no
      interfaces:
        - eno4
    br1:
      dhcp4: no
      dhcp6: no
      interfaces:
        - eno3

root@bs2020:~# nano /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg
network: {config: disabled}
root@bs2020:~# netplan apply

```

So you have to create the scripts until they fix this.

```

root@bs2020:~# nano /etc/systemd/network/br0.network
[Match]
Name=br0

[Network]
LinkLocalAddressing=no
IPv6AcceptRA=no

root@bs2020:~# nano /etc/systemd/network/br1.network
[Match]
Name=br1

[Network]
LinkLocalAddressing=no
IPv6AcceptRA=no

```

<https://bugs.launchpad.net/ubuntu/+source/nplan/+bug/1736975> <http://djanotes.blogspot.com/2018/04/anonymous-bridges-in-netplan.html>

Freaking Cloud init

Need to figure out how much damage is done here...

STARTING WITH THE HOSTNAME.

The hostname is now handled by a new command and /etc/cloud/cloud.config needs to be modified to preserve the hostname across boots.

```

feurig@bs2020:~$ sudo bash
[sudo] password for feurig:

```

```

root@bs2020:~# hostnamectl set-hostname bs2020
root@bs2020:~# nano /etc/cloud/cloud.cfg
....
# This will cause the set+update hostname module to not operate (if true)
preserve_hostname: true
...
root@bs2020:~# reboot

```

INSTALL THE ROOT USERS .

One would like for the installer to give you some options for installing the admin team but we just paste the hash from one of the other machines into the shadow password file and copy the home directories for their ssh keys. see [wiki:kb2018InstallBashHistory](#)

```
( .. tired of winning .... write up later... )
```

INSTALL ZFS

```
root@bs2020:~# apt-get install nfs-kernel-server samba-common-bin zfsutils-linux
```

- create zfs pools using lxd init.
- make servers available to each other.
- configure outgoing mail.
- install apticron

5.70.3 Apache2

Big leap in apache version. Lots of configuration changes.

Link Dump

- <https://netplan.io/examples>
- <https://websiteforstudents.com/configure-static-ip-addresses-on-ubuntu-18-04-beta/>
- <https://askubuntu.com/questions/1054350/netplan-bridge-for-kvm-on-ubuntu-server-18-04-with-static-ips> <https://stackoverflow.com/questions/33377916/migrating-lxc-to-lxd>

5.71 HOLY FUCKING AWESOME!!!!

Watch while I add a fresh disk as a mirror, resilver the pool and remove and repartition the original disk while the container using the pool is still running!!! ... make this into a structured document ...

```

root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4dev      ONLINE       0     0     0
      sdd1       ONLINE       0     0     0
      sdf        ONLINE       0     0     0
      sde        ONLINE       0     0     0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: scrub repaired 0 in 0h2m with 0 errors on Sun Aug 12 00:27:02 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4infra    ONLINE       0     0     0
      sdal       ONLINE       0     0     0

errors: No known data errors
root@bs2020:~# zpool add -n lxd4infra mirror sdb
invalid vdev specification: mirror requires at least 2 devices
root@bs2020:~# zpool add -n lxd4infra mirror sdal sdb
invalid vdev specification
use '-f' to override the following errors:
/dev/sdal is part of active pool 'lxd4infra'
/dev/sdb does not contain an EFI label but it may contain partition
information in the MBR.
root@bs2020:~# mklabel GPT /dev/sdb
bash: mklabel: command not found
root@bs2020:~# parted /dev/sdb
bash: parted: command not found
root@bs2020:~# gparted /dev/sdb
bash: gparted: command not found
root@bs2020:~# zpool add -nf lxd4infra mirror sdal sdb
invalid vdev specification
the following errors must be manually repaired:
/dev/sdal is part of active pool 'lxd4infra'
root@bs2020:~# zpool add -nf lxd4infra mirror sdb sdal
invalid vdev specification
the following errors must be manually repaired:
/dev/sdal is part of active pool 'lxd4infra'
root@bs2020:~# zpool add -nf lxd4infra mirror sdb
invalid vdev specification: mirror requires at least 2 devices
root@bs2020:~# zpool add -nf lxd4infra sdal mirror sdb
invalid vdev specification: mirror requires at least 2 devices
root@bs2020:~# zpool attach -n sdal sdb
invalid option 'n'
usage:
  attach [-f] [-o property=value] <pool> <device> <new-device>
root@bs2020:~# zpool attach sdal sdb
missing <new_device> specification
usage:
  attach [-f] [-o property=value] <pool> <device> <new-device>
root@bs2020:~# zpool attach lxd4infra sdal sdb
invalid vdev specification
use '-f' to override the following errors:
/dev/sdb does not contain an EFI label but it may contain partition
information in the MBR.
root@bs2020:~# gparted
bash: gparted: command not found
root@bs2020:~# parted
bash: parted: command not found
root@bs2020:~# apt-get install parted
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  libparted2
Suggested packages:
  libparted-dev libparted-il18n parted-doc
The following NEW packages will be installed:
  libparted2 parted
0 upgraded, 2 newly installed, 0 to remove and 0 not upgraded.
Need to get 158 kB of archives.
After this operation, 520 kB of additional disk space will be used.
Do you want to continue? [Y/n]
Get:1 http://us.archive.ubuntu.com/ubuntu xenial-updates/main amd64 libparted2 amd64 3.2-15ubuntu0.1 [115 kB]
Get:2 http://us.archive.ubuntu.com/ubuntu xenial-updates/main amd64 parted amd64 3.2-15ubuntu0.1 [42.4 kB]

```

```

Fetched 158 kB in 0s (277 kB/s)
Selecting previously unselected package libparted2:amd64.
(Reading database ... 32152 files and directories currently installed.)
Preparing to unpack .../libparted2_3.2-15ubuntu0.1_amd64.deb ...
Unpacking libparted2:amd64 (3.2-15ubuntu0.1) ...
Selecting previously unselected package parted.
Preparing to unpack .../parted_3.2-15ubuntu0.1_amd64.deb ...
Unpacking parted (3.2-15ubuntu0.1) ...
Processing triggers for libc-bin (2.23-0ubuntu10) ...
Processing triggers for man-db (2.7.5-1) ...
Setting up libparted2:amd64 (3.2-15ubuntu0.1) ...
Setting up parted (3.2-15ubuntu0.1) ...
Processing triggers for libc-bin (2.23-0ubuntu10) ...
root@bs2020:~# parted /dev/sdb
GNU Parted 3.2
Using /dev/sdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mklabel GPT
(parted) w
      align-check TYPE N           check partition N for TYPE(min|opt) alignment
      help [COMMAND]              print general help, or help on COMMAND
      mklabel,mktable LABEL-TYPE  create a new disklabel (partition table)
      mkpart PART-TYPE [FS-TYPE]  make a partition
      name NUMBER NAME             name partition NUMBER as NAME
      print [devices|free|list,all|NUMBER] display the partition table, available devices, free space, all found partitions, or a particular partition
      quit                          exit program
      rescue START END             rescue a lost partition near START and END
      resizepart NUMBER END        resize partition NUMBER
      rm NUMBER                    delete partition NUMBER
      select DEVICE                choose the device to edit
      disk_set FLAG STATE          change the FLAG on selected device
      disk_toggle [FLAG]           toggle the state of FLAG on selected device
      set NUMBER FLAG STATE        change the FLAG on partition NUMBER
      toggle [NUMBER [FLAG]]       toggle the state of FLAG on partition NUMBER
      unit UNIT                    set the default unit to UNIT
      version                      display the version number and copyright information of GNU Parted
(parted) q
Information: You may need to update /etc/fstab.

root@bs2020:~# zpool attach lxd4infra sda1 sdb
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

      NAME      STATE    READ WRITE CKSUM
      lxd4dev    ONLINE    0   0   0
        sdd1     ONLINE    0   0   0
        sdf      ONLINE    0   0   0
        sde      ONLINE    0   0   0

errors: No known data errors

pool: lxd4infra
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
       continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Tue Sep  4 09:05:14 2018
      182M scanned out of 5.38G at 10.7M/s, 0h8m to go
      181M resilvered, 3.30% done
config:

      NAME      STATE    READ WRITE CKSUM
      lxd4infra  ONLINE    0   0   0
        mirror-0  ONLINE    0   0   0
          sda1    ONLINE    0   0   0
          sdb     ONLINE    0   0   0 (resilvering)

errors: No known data errors
root@bs2020:~# packet_write_wait: Connection to 198.202.31.242: Broken pipe
steve:~ don$ ssh feurig@bs2020.suspectdevices.com
Welcome to Ubuntu 16.04.5 LTS (GNU/Linux 4.4.0-96-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

0 packages can be updated.
0 updates are security updates.

New release '18.04.1 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

Last login: Tue Sep  4 08:26:28 2018 from 75.164.203.77
feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# packet_write_wait: Connection to 198.202.31.242: Broken pipe
steve:~ don$ ssh feurig@bs2020.suspectdevices.com
Welcome to Ubuntu 16.04.5 LTS (GNU/Linux 4.4.0-96-generic x86_64)

```

```

* Documentation:  https://help.ubuntu.com
* Management:    https://landscape.canonical.com
* Support:       https://ubuntu.com/advantage

0 packages can be updated.
0 updates are security updates.

New release '18.04.1 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

Last login: Wed Sep  5 16:10:53 2018 from 75.164.203.77
feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# packet_write_wait: Connection to 198.202.31.242: Broken pipe
steve:~ don$
steve:~ don$ ssh feurig@bs2020.suspectdevices.com
Welcome to Ubuntu 16.04.5 LTS (GNU/Linux 4.4.0-96-generic x86_64)

* Documentation:  https://help.ubuntu.com
* Management:    https://landscape.canonical.com
* Support:       https://ubuntu.com/advantage

0 packages can be updated.
0 updates are security updates.

New release '18.04.1 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

Last login: Wed Sep  5 18:56:14 2018 from 75.164.203.77
feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4dev      ONLINE       0   0   0
      sdd1      ONLINE       0   0   0
      sdf       ONLINE       0   0   0
      sde       ONLINE       0   0   0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: resilvered 5.38G in 0h6m with 0 errors on Tue Sep  4 09:11:31 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4infra    ONLINE       0   0   0
      mirror-0   ONLINE       0   0   0
        sda1     ONLINE       0   0   0
        sdb      ONLINE       0   0   0

errors: No known data errors
root@bs2020:~# zpool detach -n lxd4infra sda1
invalid option 'n'
usage:
    detach <pool> <device>
root@bs2020:~# zpool detach lxd4infra sda1
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4dev      ONLINE       0   0   0
      sdd1      ONLINE       0   0   0
      sdf       ONLINE       0   0   0
      sde       ONLINE       0   0   0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: resilvered 5.38G in 0h6m with 0 errors on Tue Sep  4 09:11:31 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4infra    ONLINE       0   0   0
      sdb      ONLINE       0   0   0

errors: No known data errors
root@bs2020:~# gparted /dev/sda
bash: gparted: command not found
root@bs2020:~# parted /dev/sda
GNU Parted 3.2
Using /dev/sda
Welcome to GNU Parted! Type 'help' to view a list of commands.

```

```
(parted) mklabel GPT
Warning: The existing disk label on /dev/sda will be destroyed and all data on this disk will be lost. Do you want to continue?
Yes/No? y
(parted) q
Information: You may need to update /etc/fstab.
```

```
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4dev	ONLINE	0	0	0
sdd1	ONLINE	0	0	0
sdf	ONLINE	0	0	0
sde	ONLINE	0	0	0

```
errors: No known data errors
```

```
pool: lxd4infra
state: ONLINE
scan: resilvered 5.38G in 0h6m with 0 errors on Tue Sep 4 09:11:31 2018
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4infra	ONLINE	0	0	0
sdb	ONLINE	0	0	0

```
errors: No known data errors
```

```
root@bs2020:~# zpool attach lxd4infra sdb sda
```

```
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4dev	ONLINE	0	0	0
sdd1	ONLINE	0	0	0
sdf	ONLINE	0	0	0
sde	ONLINE	0	0	0

```
errors: No known data errors
```

```
pool: lxd4infra
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Thu Sep 6 09:24:09 2018
69.8M scanned out of 5.42G at 5.37M/s, 0h17m to go
67.9M resilvered, 1.26% done
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4infra	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
sdb	ONLINE	0	0	0
sda	ONLINE	0	0	0 (resilvering)

```
errors: No known data errors
```

```
root@bs2020:~#
```

5.72 OpenVPN on LEDE (Fail)

Now that we have a recent version of the operating system OpenVPN seems to work as advertised. Following the instructions at <https://lede-project.org/docs/user-guide/openvpn.server>. Much of the heavy lifting is done by easyRSA and MakeOpenVPN.sh.

The client setups fail if you use an empty passphrase which is good. OTOH In my initial attempts I could not get the server certificates to work with one. When in doubt read the documentation sections on the old openWRT site. It provides a little more depth but there still are some missing pieces that require more exploration (https://wiki.openwrt.org/doc/howto/vpn.openvpn#tab__using_openssl_commands_most_secure).

For the client I used tunnelblick which works well and takes the .ovpn configuration files created by this process.

Sample Install

- follow the bouncing prompt using lede user guide.

```
root@mullein:~# opkg update && opkg install openvpn-openssl openvpn-easy-rsa luci-app-openvpn
Downloading .....note additional dependencies..... Configuring kmod-tun. Configuring zlib. Configuring libopenssl. Configuring openssl-util. Configuring liblzo. Configuring openvpn-openssl. Configuring openvpn-easy-rsa. Configuring luci-app-openvpn.
root@mullein:~# cd /etc/easy-rsa root@mullein:/etc/easy-rsa# source vars NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys root@mullein:/etc/easy-rsa# clean-all NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys root@mullein:/etc/easy-rsa# build-ca NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys Generating a 2048 bit RSA private key .....+ + + .....+ + + writing new private key to 'ca.key'
```

You are about to be asked to enter information that will be incorporated into your certificate request. What you are about to enter is what is called a Distinguished Name or a DN. There are quite a few fields but you can leave some blank For some fields there will be a default value, If you enter '.', the field will be left blank.

```
Country Name (2 letter code) [US]: State or Province Name (full name) [CA]:OR Locality Name (eg, city) [SanFrancisco]:Portland Organization Name (eg, company) [Fort-Funston]:SuspectDevices Organizational Unit Name (eg, section) [MyOrganizationalUnit]:3dAngst Common Name (eg, your name or your server's hostname) [Fort-Funston CA]:mullein Name [EasyRSA]:mullein Email Address [me@myhost.mydomain]:don@suspectdevices.com
```

- plan on the next step taking so long you will probably have to reconnect and pick up where you were...

```
root@mullein:/etc/easy-rsa# build-dh NOTE: If you run ./clean-all, I will be doing a rm -rf on /etc/easy-rsa/keys Generating DH parameters, 2048 bit long safe prime, generator 2 This is going to take a long time .... They are not kidding
.....+ + + +
```

- continue to follow the bouncing prompt

```
root@mullein:/etc/easy-rsa# build-key-server mullein ..... answer the questions .... A challenge password []: An optional company name []: Using configuration from /etc/easy-rsa/openssl-1.0.0.cnf Check that the request matches the signature Signature ok The Subject's Distinguished Name is as follows countryName :PRINTABLE:'US' stateOrProvinceName :PRINTABLE:'OR' localityName :PRINTABLE:'Portland' organizationName :PRINTABLE:'SuspectDevices' organizationalUnitName:PRINTABLE:'3dAngst' commonName :PRINTABLE:'mullein' name :PRINTABLE:'mullein' emailAddress :IA5STRING:'don@suspectdevices.com' Certificate is to be certified until Oct 23 23:46:35 2027 GMT (3650 days) Sign the certificate? [y/n]:y 1 out of 1 certificate requests certified, commit? [y/n]:y Write out database with 1 new entries Data Base Updated root@mullein:/etc/easy-rsa# openvpn --genkey --secret /etc/easy-rsa/keys/ta.key
```

- set up the network and firewall rules.

```
root@mullein:/etc/easy-rsa# openvpn --genkey --secret /etc/easy-rsa/keys/ta.key root@mullein:/etc/easy-rsa# uci set network.vpn0="interface" root@mullein:/etc/easy-rsa# uci set network.vpn0.ifname="tun0" root@mullein:/etc/easy-rsa# uci set network.vpn0.proto="none" root@mullein:/etc/easy-rsa# uci set network.vpn0.auto="1" root@mullein:/etc/easy-rsa# uci commit network root@mullein:/etc/easy-rsa# uci add firewall rule cfg1892bd root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].name="Allow-OpenVPN-Inbound" root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].target="ACCEPT"
```



```

root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].src="wan" root@mullein:/etc/easy-rsa# uci set
firewall.@rule[-1].proto="udp" root@mullein:/etc/easy-rsa# uci set firewall.@rule[-1].dest_port="1194" root@mullein:/etc/
easy-rsa# uci add firewall zone cfg19dc81 root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].name="vpn" root@mullein:/
etc/easy-rsa# uci set firewall.@zone[-1].input="ACCEPT" root@mullein:/etc/easy-rsa# uci set
firewall.@zone[-1].forward="ACCEPT" root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].output="ACCEPT"
root@mullein:/etc/easy-rsa# uci set firewall.@zone[-1].masq="1" root@mullein:/etc/easy-rsa# uci set
firewall.@zone[-1].network="vpn0" root@mullein:/etc/easy-rsa# uci add firewall forwarding cfg1aad58 root@mullein:/etc/
easy-rsa# uci set firewall.@forwarding[-1].src="vpn" root@mullein:/etc/easy-rsa# uci set
firewall.@forwarding[-1].dest="wan" root@mullein:/etc/easy-rsa# uci add firewall forwarding cfg1bad58 root@mullein:/etc/
easy-rsa# uci set firewall.@forwarding[-1].src="vpn" root@mullein:/etc/easy-rsa# uci set firewall.@forwarding[-1].dest="lan"
root@mullein:/etc/easy-rsa# uci commit firewall root@mullein:/etc/easy-rsa# /etc/init.d/network reload .... root@mullein:/etc/
easy-rsa# /etc/init.d/firewall reload ....

```

- check ip forwarding

```

root@mullein:/etc/easy-rsa# cat /proc/sys/net/ipv4/ip_forward 1

```

- edit /etc/config/openvpn, enable and restart daemon.

```

root@mullein:/etc/easy-rsa# nano /etc/config/openvpn ... add the following (change name, cert, and key to match your server)
...

```

5.73 <https://lede-project.org/docs/user-guide/openvpn.server>

```

config openvpn 'mullein' option enabled '1' option dev 'tun' option port '1194' option proto 'udp' option status '/var/log/
openvpn_status.log' option log '/tmp/openvpn.log' option verb '3' option mute '5' option keepalive '10 120' option persist_key
'1' option persist_tun '1' option user 'nobody' option group 'nogroup' option ca '/etc/easy-rsa/keys/ca.crt' option cert '/etc/easy-
rsa/keys/mullein.crt' option key '/etc/easy-rsa/keys/mullein.key' option dh '/etc/easy-rsa/keys/dh2048.pem' option mode 'server'
option tls_server '1' option tls_auth '/etc/easy-rsa/keys/ta.key 0' option server '10.9.0.0 255.255.255.0' option topology
'subnet' option route_gateway 'dhcp' option client_to_client '1' list push 'persist-key' list push 'persist-tun' list push 'redirect-
gateway def1' # allow your clients to access to your network list push 'route 192.168.2.0 255.255.255.0' # push DNS to your
clients list push 'dhcp-option DNS 192.168.2.1' option comp_lzo 'no'

```

```

root@mullein:/etc/easy-rsa# /etc/init.d/openvpn start root@mullein:/etc/easy-rsa# /etc/init.d/openvpn enable root@mullein:/
etc/easy-rsa# cat /tmp/openvpn.log ... Thu Oct 26 00:22:46 2017 OpenVPN 2.4.3 mipsel-openwrt-linux-gnu [SSL (OpenSSL)]
[LZO] [LZ4] [EPOLL] [MH/PKTINFO] [AEAD] .... Thu Oct 26 00:22:46 2017 MULTI: multi_init called, r=256 v=256 Thu Oct
26 00:22:46 2017 IFCONFIG POOL: base=10.9.0.2 size=252, ipv6=0 Thu Oct 26 00:22:46 2017 Initialization Sequence
Completed ...

```

- create client cert.

```

root@mullein:~# cd /etc/easy-rsa/ root@mullein:/etc/easy-rsa# source vars NOTE: If you run ./clean-all, I will be doing a rm -rf
on /etc/easy-rsa/keys root@mullein:/etc/easy-rsa# build-key-pkcs12 donathome ... writing new private key to 'donathome.key'
.... Country Name (2 letter code) [US]: State or Province Name (full name) [CA]:OR Locality Name (eg, city)
[SanFrancisco]:Portland Organization Name (eg, company) [Fort-Funston]:SuspectDevices Organizational Unit Name (eg,
section) [MyOrganizationalUnit]:3dAngst Common Name (eg, your name or your server's hostname) [donathome]:viscious
Name [EasyRSA]:DonAtHome Email Address [me@myhost.mydomain]:don@suspectdevices.com

```

```

Please enter the following 'extra' attributes to be sent with your certificate request A challenge password []:XXXXXXXXXXXX
An optional company name []:Its Late ... Certificate is to be certified until Oct 24 02:49:46 2027 GMT (3650 days) Sign the
certificate? [y/n]:y

```

```

1 out of 1 certificate requests certified, commit? [y/n]y Write out database with 1 new entries Data Base Updated Enter
Export Password: Verifying - Enter Export Password: root@mullein:/etc/easy-rsa# openssl rsa -in /etc/easy-rsa/keys/
donathome.key -des3 -out /etc/easy-rsa/keys/donathome.3des.key writing RSA key Enter PEM pass phrase: Verifying - Enter
PEM pass phrase: root@mullein:/etc/easy-rsa#

```

- MakeOpenVPN.sh script (install missing dependencies)

```

root@mullein:/etc/easy-rsa# cd keys root@mullein:/etc/easy-rsa/keys# wget https://gist.github.com/ivanmarban/
57561e2bacf3b3a709426d353d2b6584/raw/30bf3c86fbc95a0a5d53d0aac348bcebd9aa2eb/MakeOpenVPN.sh -O /etc/
easy-rsa/keys/MakeOpenVPN.sh wget: SSL support not available, please install one of the libstream-ssl libraries as well as the

```

```
ca-bundle and ca-certificates packages. root@mullein:/etc/easy-rsa/keys# opkg update && opkg install libstream-openssl ca-
certificates ... root@mullein:/etc/easy-rsa/keys# wget https://gist.githubusercontent.com/ivanmarban/
57561e2bacf3b3a709426d353d2b6584/raw/30bf3c86fbc95a0a5d53d0aac348bcebd9aa2eb/MakeOpenVPN.sh -O /etc/
easy-rsa/keys/MakeOpenVPN.sh Downloading 'https://gist.githubusercontent.com/ivanmarban/
57561e2bacf3b3a709426d353d2b6584/raw/30bf3c86fbc95a0a5d53d0aac348bcebd9aa2eb/MakeOpenVPN.sh'
Connecting to 151.101.52.133:443 Writing to '/etc/easy-rsa/keys/MakeOpenVPN.sh' /etc/easy-rsa/keys/M 100% |*****| 1839
0:00:00 ETA Download completed (1839 bytes) root@mullein:/etc/easy-rsa/keys# chmod oug+x MakeOpenVPN.sh
```

- configure and run script.

```
root@mullein:/etc/easy-rsa/keys# nano Default.txt ... Add the following, Adjust host name accordingly .... client dev tun proto
udp remote mullein.suspectdevices.com 1194 resolv-retry infinite nobind mute-replay-warnings ns-cert-type server key-
direction 1 verb 1 mute 20 comp-lzo no root@mullein:/etc/easy-rsa/keys# ./MakeOpenVPN.sh Please enter an existing Client
Name: donathome Client's cert found: donathome Client's Private Key found: donathome.3des.key CA public Key found:
ca.crt tls-auth Private Key found: ta.key Done! donathome.ovpn Successfully Created. root@mullein:/etc/easy-rsa/keys# ls
01.pem ca.crt donathome.key index.txt.old mullein.key myvpn.key 02.pem ca.key donathome.ovpn knight.crt mullien.crt
serial 03.pem dh2048.pem donathome.p12 knight.csr mullien.csr serial.old 04.pem donathome.3des.key index.txt knight.key
mullien.key ta.key Default.txt donathome.crt index.txt.attr mullein.crt myvpn.crt MakeOpenVPN.sh donathome.csr
index.txt.attr.old mullein.csr myvpn.csr root@mullein:/etc/easy-rsa/keys# ./MakeOpenVPN.sh Please enter an existing Client
Name: donathome Client's cert found: donathome Client's Private Key found: donathome.3des.key CA public Key found:
ca.crt tls-auth Private Key found: ta.key Done! donathome.ovpn Successfully Created.
```

References (Link Dump)

- <https://help.my-private-network.co.uk/support/solutions/articles/24000005597-openwrt-lede-openvpn-setup>
- https://lede-project.org/docs/user-guide/openvpn.server#setup_clients
- <https://steemit.com/openwrt/@rbrthnk/vpn-pptp-router-with-openwrt-lede-tutorial-super-easy>
- https://lede-project.org/docs/user-guide/tunneling_interface_protocols
- https://www.softether.org/4-docs/2-howto/9.L2TP/IPsec_Setup_Guide_for_SoftEther_VPN_Server
- https://wiki.gentoo.org/wiki/IPsec_L2TP_VPN_server
- http://connect.rbhs.rutgers.edu/vpn/Mac_OSX_Native_VPN_Client_Overview.pdf
- <http://cookbook.fortinet.com/ipsec-vpn-native-mac-os-client-54/>
- <https://www.howtogeek.com/216209/how-to-connect-your-mac-to-any-vpn-and-automatically-reconnect/>
- <https://tunnelblick.net/cInstall.html>
- <https://forum.lede-project.org/t/configuring-lede-router-with-a-pppoe-modem-router/5348/2>
- <https://wiki.openwrt.org/doc/howto/openconnect-setup>
- https://wiki.gavowen.ninja/doku.php?id=lede:openconnect#tab__pki_templates
- <https://lede-project.org/docs/user-guide/openvpn.server>
- https://wiki.openwrt.org/doc/howto/vpn.openvpn#tab__traditional_tun_client

5.74 OpenWRT Notes

At a very minimum open the ssh port so that the router can be managed from the outside. Then disable logins (ssh keys only) in `/etc/dropbear`.

```
root@OpenWrt:/etc/config# opkg update
root@OpenWrt:/etc/config# opkg install nano

root@OpenWrt:/etc/config# nano /etc/config/firewall
```

add the following

```
config redirect
    option target 'DNAT'
    option src 'wan'
    option dest 'lan'
    option proto 'tcp'
    option dest_ip '192.168.1.1'
    option dest_port '22'
    option name 'sshplease'
    option src_dport '2222'
```

5.74.1 allowing access to dell IDRAC 6 and server forward

5.74.2 firewall setup on vpn

In order to get at the idrac and access BS2020 via ssh the following rules were added to `/etc/config/firewall`

```
config redirect
    option target 'DNAT'
    option src 'wan'
    option dest 'lan'
    option proto 'tcp'
    option dest_ip '192.168.1.158'
    option dest_port '22'
    option name 'sshtobernie'
    option src_dport '22'

# idrac 6 redirections
config redirect
    option target 'DNAT'
    option src 'wan'
    option dest 'lan'
    option proto 'tcp'
    option dest_ip '192.168.1.121'
    option dest_port '443'
    option name 'idracplease1'
    option src_dport '443'

config redirect
    option target 'DNAT'
    option src 'wan'
    option dest 'lan'
    option proto 'tcp'
    option dest_ip '192.168.1.121'
    option dest_port '4433'
    option name 'idracplease2'
    option src_dport '4433'

config redirect
    option target 'DNAT'
    option src 'wan'
    option dest 'lan'
    option proto 'tcp'
    option dest_ip '192.168.1.121'
    option dest_port '443'
    option name 'idracplease3'
    option src_dport '443'

config redirect
    option target 'DNAT'
    option src 'wan'
    option dest 'lan'
    option proto 'tcp'
    option dest_ip '192.168.1.121'
    option dest_port '623'
    option name 'idracplease4'
    option src_dport '623'
```

Just to be paranoid we "`#uci show`" to make sure UCI picks up the rules then we "`#uci commit`" and reboot the router.

at this point we have full access to the servers idrac6

5.74.3 Related Pages

OpenVPN attempt #2

[wiki:OpenVPNOnLEDE OpenVPN on LEDE]

Adventures in deploying OpenWRT/LEDE

- [wiki:OpenWRTonMR3020 Open WRT on TP-Link MR3020]
- [wiki:OpenWRTonLinkSysEA3500 Open WRT on LinkSYS EA3500]

5.75 LEDE on EA3500

This guy required me to upload the 15.04 and sys upgrade. Otherwise not a huge deal.

- <https://wiki.openwrt.org/toh/linksys/ea3500>

5.76 Building old LEDE firmware

```
root@sandbox:/home/openwrt/15.05.1/OpenWrt-ImageBuilder-15.05.1-ar71xx-generic.Linux-x86_64/bin/ar71xx# make info
....
root@sandbox:/home/openwrt/15.05.1/OpenWrt-ImageBuilder-15.05.1-ar71xx-generic.Linux-x86_64/bin/ar71xx# make image PROFILE=TLMR3020 PACKAGES="nano"
....
```

getting firmware onto local system. the stock firmware will not accept a firmware that is not the same name as a stock firmware.

```
viscious:vpn don$ scp feurig@sandbox:/home/openwrt/15.05.1/OpenWrt-ImageBuilder-15.05.1-ar71xx-generic.Linux-x86_64/bin/ar71xx/openwrt-15.05.1-ar71xx-generic-tl-mr3020-v1-squashfs-factory.bin .
viscious:vpn don$ mv openwrt-15.05.1-ar71xx-generic-tl-mr3020-v1-squashfs-factory.bin mr3020nv1_en_3_17_2_up_boot(150921).bin
```

At this point you can telnet to the router and reset the root password (which will disable telnet and enable ssh)

related

- [wiki:LEDE LEDE]

References

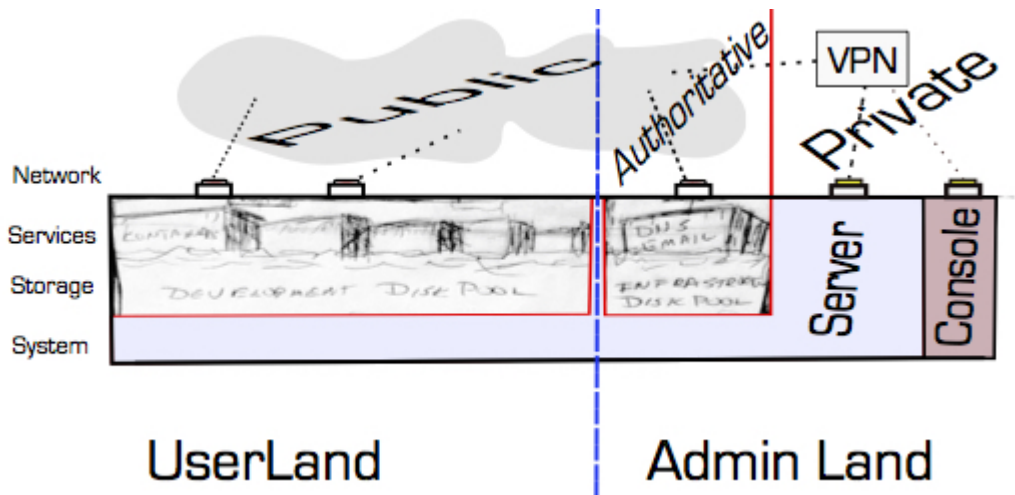
- <https://nicolas314.wordpress.com/2015/12/09/openwrt-on-mr3020/>
- <https://wolfgang.reutz.at/2012/04/12/openwrt-on-tp-link-mr3020-as-infopoint-with-local-webserver/>
- <https://blog.philippklaus.de/2012/03/openwrt-on-a-tp-link-tl-mr3020-router/>
- <https://openwrt.org/docs/guide-user/additional-software/imagebuilder>

5.77 OpenWRT E900 Firmware Build

```
feurig@sandbox:~$ cd /home/openwrt/current/openwrt-imagebuilder-18.06.1-brcm47xx-mips74k.Linux-x86_64/
feurig@sandbox:/ho...64$ sudo cat ~joe/.ssh/authorized_keys ~feurig/.ssh/authorized_keys >files/etc/dropbear/authorized_keys
feurig@sandbox:/ho...64$ make image PROFILE=linksys-e900-v1 PACKAGES="nano sudo shadow shadow-utils shadow-vipw -luci -ppp -ppp-mod-pppoe -odhcp6c -odhcpd-ipv6only"
FILES="files/"
....
Calculating checksums...
feurig@sandbox:/ho...64$ ls bin/targets/brcm47xx/mips74k/
openwrt-brcm47xx-mips74k-asus-rt-ac53u-squashfs.trx
....
brcm47xx-mips74k-linksys-e900-v1-squashfs.bin
...
openwrt-brcm47xx-mips74k-linksys-e2500-v2.1-squashfs.bin      ...
feurig@sandbox:/home/openwrt/current/openwrt-imagebuilder-18.06.1-brcm47xx-mips74k.Linux-x86_64$
```

5.78 Server Modernization

5.78.1 Overview



Phase I

Phase one of the server modernization shifted away from multipurposed servers and kvms to lxc/lxd based containers.

- Moving all legacy system functions onto separate linux containers isolated from each other.
- Use mirrored disk systems to insure that disk corruption does not lead to data corruption.
- Start giving a shit about the systems, code, and sites on them.
- Own your code/data. (If your free code hosting system is shutdown or taken over by Microsoft is it really free)

Server Modernization Phase II

Phase two extends on this by integrate Ansible into system maintenance tasks.

- Integrate Ansible into system maintenance tasks
- Reevaluate Centos and other RPM based containers built using playbooks vs profiles/scripts/cloud-init *while maintaining current security model*
- Develop off site backup strategy.
- Clean up the cruft (If it doesn't bring you joy DTMFA)

SMP III Make Shit Happen / Own Your Shit

- Work on secure and efficient traffic in and out of home lans (Privoxy,DNS based ad blocking,squid etc)
- Continue to refine server operation/maintanance.
- Build Gitlab and other alternatives to trac/git and evaluate workflows.
- Deploy off site backup strategy.
- Build out content.
- Start new projects.
- Distribute data and backups over the network to home servers.
- [Document home server/network setup](#)

Goals.

- Security
- Flexibility
- Simplification

Isolation

- network
- performance
- disk

5.78.2 Hardware

At present the environment contains a vpn capable router (Knight) and two enterprise class servers

- bs2020 , a Dell PowerEdge R610 [[br]]and
- kb2018 a HP ProLiant DL380 (g7) .

5.78.3 Network


The network is divided into 3 segments

- 192.168.31.0/24 a private administrative lan
- tbd.tbd.tbd.tbd/? a private vpn for home offices
- 198.202.31.129/25 A public facing lan.

The hosts themselves do not have any public facing interfaces and are only accessible through the admin lan. The containers which handle all public facing work do so via an anonymous bridge configuration, allowing them to access the internet directly without allowing external access to the servers.

bs2020 ports				
port	Interface	IP Address/mask	linux device	purpose
1	eno1	192.168.31.158/24	eno1	internal / admin lan
2	?	???.?/??	eno2	vpn for home/office networks
3	br1	0.0.0.0/0	eno3	Public Interface for infrastructure servers
4	br0	0.0.0.0/0	eno4	Public Interface for dev/deployment servers
idrac		192.168.31.121/24		remote console


As Drawn	As Deployed.
----------	--------------



				kb2018 ports
port	Interface	IP Address/mask	linux device	purpose
4	enp4s0f1	192.168.31.159/24	enp4s0f1	internal / admin lan
3	enp4s0f0	???.?/?	enp4s0f0	vpn for home/office networks
2	br1	0.0.0.0/32	enp3s0f1	Public Interface for infrastructure servers
1	br0	0.0.0.0/32	enp3s0f0	Public Interface for dev/deployment servers
ilo		192.168.31.119/24		remote console

As Drawn

As Dep



See: <https://bitbucket.org/suspectdevicesadmin/ansible/src/master/hosts> which is built referencing a google doc with proposed allocations

5.78.4 Server OS, Filesystems and Disk layout

The servers are both running a standard install Ubuntu Server LTS, along with the Canonical supported LXD "snap". Outside of zfs not much is added to the stock installation. This is intentional. Since the real work is done by the containers the host os is considered disposable and can be rebuilt without effecting production.

Disk Layout

The system disks on both servers use hardware raid 1+0 mirroring. The containers are able to take advantage of zfs mirroring and caching.

bs2020 disks					
disk	device/pool	bay	type	mount point(s)	purpose/notes
Host Machine Disks					
sdg	/dev/sdg	0	ext4	/	root filesystem (hardware raid)
sdg	/dev/sdg	1	ext4	/	mirror
sda1	/dev/sda1	external	ext4	/archive	backup staging
development zfs pool					
sdc	devel	2	zfs	/var/lib/lxd/storage-pools/devel	dev/deployment (www,trac,usw)
sdd	devel	3	zfs		mirror
development zfs pool					
sdd	infra	4	zfs	/var/lib/lxd/storage-pools/infra	infrastructure (email,dns,usw)
sde	infra	5	zfs		mirror

On kb2018 the second pair of disks are Solid State. The first partition on each is a mirrored pair for the infrastructure zfs pool. The remaining partitions are for zfs caching.

kb2018 disks					
disk	device/pool	bay	type	mount point(s)	purpose/notes
Host Machine Disks					
sda	/dev/sda	0	ext4	/	root filesystem (hardware raid)
sda	/dev/sda	1	ext4	/	mirror
infrastructure zfs pool					
sdb1	infra	2	zfs	/var/lib/lxd/storage-pools/infra	infrastructure (email,dns,usw)
sdc1	infra	3	zfs		mirror
development zfs pool					
sdd	devel	4	zfs	/var/lib/lxd/storage-pools/devel	dev/deployment (www,trac,usw)
sde	devel	5	zfs		mirror
sdb2	devel	2	zfs		zfs cache (proposed)

Hardware raid on the DL380

The raid controller on the Dell allows a mixing of hardware raid and direct hot swappable connections. The HP 420i does only hardware raid or direct connections (HBA) but not both. Since we use the hardware raid the remaining disks need to be configured using the sscli or the raid controllers bios. See: [DudeWhereAreMyDisks](#)

5.78.5 Containers

Work previously done by standalone servers is now done through LXD managed containers. [\[#fn1 \(1\)\]](#) An up to date list of containers is maintained at <https://bitbucket.org/suspectdevicesadmin/ansible/src/master/hosts>

5.78.6 Ansible

Ansible is used to make most tasks reasonable including. * creating containers * updating containers * updating admin passwords and ssh keys. * accessing

5.79 Tasks: Accessing Hosts

bs2020/kb2020 ssh access

The host machines for the containers can be accessed through the admin lan. Currently this is done through ssh redirection. Eventually it will require a vpn connection. Only ssh key access is allowed and root is not allowed to login. To escalate privileges requires sudo.

CURRENT SSH PORT MAPPINGS TO VPN.SUSPECTDEVICES.COM

port	destination
22	bs2020 ssh via admin lan
222	bs2020 racadm / serial console via ssh
2222	knight / vpn
22222	kb2018 ssh via admin lan
22223	kb2018 hpILO / serial console via ssh

note: as of a few updates ago you have to tell apples ssh client to use ssh-dss as below

```

steve:~ don$ ssh -p22223 -oHostKeyAlgorithms+=ssh-dss feurig@bs2020.suspectdevices.com
User:feurig logged-in to kb2018.suspectdevices.com(192.168.31.119 / FE80::9E8E:99FF:FE0C:BAD8)
iLO 3 Advanced for BladeSystem 1.88 at Jul 13 2016
Server Name: kb2018
Server Power: On

hpil0-> vsp

Virtual Serial Port Active: COM2

Starting virtual serial port.
Press 'ESC (' to return to the CLI Session.

Ubuntu 18.04.1 LTS kb2018 ttyS1

kb2018 login: <ESC> (
hpil0-> exit
steve:~ don$ ssh -p 222 feurig@vpn.suspectdevices.com
...
/admin1-> console com2
Connected to Serial Device 2. To end type: ^\

Ubuntu 18.04.1 LTS bs2020 ttyS1

bs2020 login: <CTL> \
/admin1-> exit
CLP Session terminated
Connection to vpn.suspectdevices.com closed.
steve:~ don$

```

_ if the serial port is still in use do the following _

```
Virtual Serial Port is currently in use by another session.
hpi0-> stop /system1/oemhp_vsp1
```

bs2020/kb2018 graphical console access

bs2020 allows complete control of the system via a Dell Idrac 6 controller. This also requires access to the admin lan. This is described on the [wiki:NotesOnIdrac6 Idrac 6 page] kb2020 allows similar using the on board described on the [wiki:NotesOnILO3 ILO 3 Notes page.]

ssh access to containers

The susdev profile adds ssh keys and sudo passwords for admin users allowing direct ssh access to the container.

```
steve:~ don$ ssh feurig@ian.suspectdevices.com
...
feurig@ian:~$
```

The containers can be accessed directly from the lxc/lxd host as root

```
root@bs2020:~# lxc exec harvey bash
root@harvey:~# apt-get update&&apt-get -y dist-upgrade&& apt-get -y autoremove
```

5.79.1 Updating dns

Dns is provided by bind , The zone files have been consolidated into a single directory under /etc/bind/zones on naomi (dns.suspectdevices.com).

```
root@naomi:/etc/bind/zones# nano suspectdevices.hosts
...
@           IN      SOA  dns1.digithink.com. don.digithink.com (
                2018080300 10800 3600 3600000 86400 )
;           ^^^ update ^^
; ... make some changes ...
morgan      IN      A      198.202.31.224
git         IN      CNAME  morgan
...
root@naomi:/etc/bind/zones# service bind9 reload
root@naomi:/etc/bind/zones# tail /var/log/messages
...
Sep  3 08:10:04 naomi named[178]: zone suspectdevices.com/IN: loaded serial 2018080300
Sep  3 08:10:04 naomi named[178]: zone suspectdevices.com/IN: sending notifies (serial 2018080300)
Sep  3 08:10:04 naomi named[178]: client 198.202.31.132#56120 (suspectdevices.com): transfer of 'suspectdevices.com/IN': AXFR-style IXFR started (serial 2018080300)
Sep  3 08:10:04 naomi named[178]: client 198.202.31.132#56120 (suspectdevices.com): transfer of 'suspectdevices.com/IN': AXFR-style IXFR ended
Sep  3 08:10:04 naomi named[178]: client 198.202.31.132#47381: received notify for zone 'suspectdevices.com'
```

5.79.2 Updating Hosts / Containers

When updates are available Apticron sends us an email. We prefer this to autoupdating our hosts as it helps us maintain awareness of what issues are being addressed and does not stop working when there are issues. All hosts in /etc/asnsible/hosts on kb2018 should be updated using the following add hoc command.

```
feurig@kb2018:~$ sudo bash
....
root@kb2018:~# ansible pets -m raw -a "update.sh"
```

<https://bitbucket.org/suspectdevicesadmin/ansible/src/master/files/update.sh>

5.79.3 Creating containers

```
ansible-playbook playbooks/create-lxd-containers.yml
```

https://bitbucket.org/suspectdevicesadmin/ansible/src/master/roles/create_lxd_containers/tasks/main.ymlYOU ARE HERE.....
documenting the ansible script to create containers.

5.79.4 Backing Up Containers

Backing up containers using ansible is depreciated. A python script and cron tab create nightly snapshots and moves them to bs2020.

```
cd /etc/ansible ;screen -L time ansible-playbook playbooks/backup-lxd-containers.yml -vvv -i importants
```

https://bitbucket.org/suspectdevicesadmin/ansible/src/master/roles/snapshot_lxd_containers/tasks/main.yml

5.80 links.... (tbd)

5.81 PlatformIO

I am looking to replace the Arduino framework with platform io and Xcode with Atom. The first test of this will be to program the [wiki:Esp8266] before moving back to the [wiki:Samd21 M0], [wiki:LeaflabsMaple Maple], and other [wiki:Arduino] boards.

Platformio is installed via pip.

```
root@bob2:~# apt-get install python-pip
... pip says we should upgrade ...
root@bob2:~# pip install --upgrade pip
...
root@bob2:~# pip install platformio
```

once installed you can use it to get most of its dependencies. *(not sure I like the way it stores everything int its own space in my home directory)*

```
don@bob2:~$ cd Documents
don@bob2:~/Documents$ mkdir piotest
don@bob2:~/Documents$ cd piotest
don@bob2:~/Documents/piotest$ platformio init board=thingdev
....
save current ino file to directory and move it to src
don@bob2:~/Documents/piotest$ mv mDNS_Web_Server/mDNS_Web_Server.ino src
don@bob2:~/Documents/piotest$ platformio run --target upload
....
```

5.81.1 Linkdump

- <https://www.penninkhof.com/2015/12/1610-over-the-air-esp8266-programming-using-platformio/>
- <https://blog.openenergymonitor.org/2016/06/esp8266-ota-update/>
- <https://randomnerdtutorials.com/esp8266-ota-updates-with-arduino-ide-over-the-air/>
- <https://github.com/openenergymonitor/EmonESP>
- <https://blog.squix.org/2016/06/esp8266-continuous-delivery-pipeline-push-to-production.html>
- <https://www.thingforward.io/techblog/2016-11-22-getting-started-with-platformio-and-esp8266htmlmarkdown.html>
- https://esp8266.github.io/Arduino/versions/2.0.0/doc/ota_updates/ota_updates.html
- <https://www.bakke.online/index.php/2017/06/02/self-updating-ota-firmware-for-esp8266/>

5.82 RecentChanges

5.83 Redmine Install

Redmine installation is documented at the git repo for the documentation for configuring the server. <https://github.com/feurig/redmine-configuration>

5.84 Foobarred zfs filesystem

When replacing our new disks there were hard errors on the disk being resilvered from. (spot the error....)

```

root@bs2020:~# zpool status -v
pool: devel
state: DEGRADED
status: One or more devices has experienced an error resulting in data
corruption. Applications may be affected.
action: Restore the file in question if possible. Otherwise restore the
entire pool from backup.
see: http://zfsonlinux.org/msg/ZFS-8000-8A
scan: resilvered 9.95G in 0h6m with 4 errors on Fri Nov 16 14:04:48 2018
config:

NAME                STATE      READ WRITE CKSUM
devel                DEGRADED   10    0    0
  mirror-0
    scsi-35000c50047d16807 DEGRADED   40    0   12 too many errors
    scsi-35000c50047d0926f ONLINE      0    0   27

errors: Permanent errors have been detected in the following files:

devel/containers/naomi7oct2018:/rootfs/home/feurig/mailstuff.tgz
devel/containers/naomi7oct2018:/rootfs/usr/lib/x86_64-linux-gnu/gconv/UTF-7.so
devel/containers/naomi7oct2018:/old.rootfs/home/feurig/var/lib/lxc/naomi/rootfs/home/feurig/mailstuff.tgz
devel/containers/naomi7oct2018:/old.rootfs/home/feurig/var/lib/lxc/naomi/rootfs/home/don/Maildir/.INBOX.arduino/cur/
1441292156.M90971P8887.bernie,S=8756,W=8933:2,Sab

pool: infra
state: ONLINE
scan: scrub repaired 0B in 0h0m with 0 errors on Sun Nov 11 00:24:23 2018
config:

NAME                STATE      READ WRITE CKSUM
infra                ONLINE      0    0    0
  mirror-0
    scsi-35000cca00b33a264 ONLINE      0    0    0
    scsi-350000395a8336d34 ONLINE      0    0    0

errors: No known data errors
root@bs2020:~#

sync [pool] ...
root@bs2020:~# zpool detach devel scsi-35000c50047d0926f
root@bs2020:~# zpool status
pool: devel
state: DEGRADED
status: One or more devices has experienced an error resulting in data
corruption. Applications may be affected.
action: Restore the file in question if possible. Otherwise restore the
entire pool from backup.
see: http://zfsonlinux.org/msg/ZFS-8000-8A
scan: resilvered 9.95G in 0h6m with 4 errors on Fri Nov 16 14:04:48 2018
config:

NAME                STATE      READ WRITE CKSUM
devel                DEGRADED   10    0    0
  scsi-35000c50047d16807 DEGRADED   40    0   12 too many errors

errors: 4 data errors, use '-v' for a list

pool: infra
state: ONLINE
scan: scrub repaired 0B in 0h0m with 0 errors on Sun Nov 11 00:24:23 2018
config:

NAME                STATE      READ WRITE CKSUM
infra                ONLINE      0    0    0
  mirror-0
    scsi-35000cca00b33a264 ONLINE      0    0    0
    scsi-350000395a8336d34 ONLINE      0    0    0

errors: No known data errors

pool: devel
state: DEGRADED
status: One or more devices is currently being resilvered. The pool will
continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Fri Nov 16 14:49:37 2018
  1.32G scanned out of 9.95G at 16.0M/s, 0h9m to go
  1.32G resilvered, 13.31% done
config:

NAME                STATE      READ WRITE CKSUM

```

```

devel                DEGRADED    10    0    0
replacing-0          DEGRADED    0    0    0
  scsi-35000c50047d16807 DEGRADED    40    0    12 too many errors
  scsi-35000c50047d0926f ONLINE      0    0    0 (resilvering)

errors: 4 data errors, use '-v' for a list

pool: infra
state: ONLINE
scan: scrub repaired 0B in 0h0m with 0 errors on Sun Nov 11 00:24:23 2018
config:

NAME                STATE      READ WRITE CKSUM
infra                ONLINE      0    0    0
  mirror-0           ONLINE      0    0    0
    scsi-35000cca00b33a264 ONLINE      0    0    0
    scsi-350000395a8336d34 ONLINE      0    0    0

errors: No known data errors
root@bs2020:~# zpool status[ 9307.615155] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9309.788008] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9312.115335] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9313.886154] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9315.603474] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9421.337403] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9424.263000] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9426.668087] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9428.468338] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9430.192036] print_req_error: critical medium error, dev sdc, sector 45059769
[ 9502.892589] print_req_error: critical medium error, dev sdc, sector 99975477
[ 9505.406331] print_req_error: critical medium error, dev sdc, sector 99975477
[ 9628.405254] print_req_error: critical medium error, dev sdc, sector 106690474
[ 9630.596015] print_req_error: critical medium error, dev sdc, sector 106690474
[ 9638.557126] print_req_error: critical medium error, dev sdc, sector 107074277
[ 9641.058573] print_req_error: critical medium error, dev sdc, sector 107074277

```

5.85 Start Using the F words.

Ubuntu's next long term support version 20.04 (Focal Fossa) is set to be released by the end of the month. In order to be prepared we should start using them.

5.85.1 Our experience so far

Lxc container

I ran up a few new containers using our profiles and test for status around [ticket:42 recent issues with "resolve"d].

I am still working on whether or not this is resolved or if I broke it trying to get resolv'd to listen to our servers using the init profiles.

The images are split into ubuntu/focal and ubuntu/focal/cloud. The cloud image picks up most of the profile changes and shows the most promise so far. This is a nice change given that most of the non-lts images required tweaking before just working.

upgraded containers

IAN / WORDPRESS SITE

php7.2->7.4upgrade broke the site.

Removed mods-enabled/php7.2*

```
root@ian:/etc/apache2/mods-enabled# mv php7.2.* /tmp/
```

and enabled php7.4

```
root@ian:/etc/apache2/mods-enabled# ln -s ../mods-available/php7.4* .
```

susdev20 (changes to profile)

- update joe's passed and keys. (this does not replace ticket:44)
- remove resolv'd file creation
- add pi-hole update to update script.

Lxc server

I performed a do-release-upgrade -d on Joey.

```
root@joey:~# do-release-upgrade -d -m server
```

It was flawless except that I had to run

```
#netplan apply
```

From the *console* which pretty much fucks any chance of doing Bernie next. (At least until [ticket:36 the issues with console redirection are resolved]).

BS2020

Once the console came back up I upgraded Bernie. ZFS pools needed to be updated after upgrade. Otherwise everything went pretty well.

Notes

Major Changes. * Lxd 4.0 * Php7.4 * Postgresql 12.

5.86 SuspectDevices

- Midi / MissingLink
- Midi-usb on maple bacon.

5.87 System Updates (for gihon)

When you log into your ubuntu cloud server it will greet you with most of what need to know to keep it up and running.

```
Welcome to Ubuntu 14.04.4 LTS (GNU/Linux 3.13.0-74-generic x86_64)

* Documentation:  https://help.ubuntu.com/

System information as of Thu Apr  7 19:51:52 UTC 2016

System load:  0.0           Processes:            131
Usage of /:   70.9% of 29.39GB Users logged in:      0
Memory usage: 29%          IP address for eth0: 172.31.16.108
Swap usage:   0%

Graph this data and manage this system at:
https://landscape.canonical.com/

Get cloud support with Ubuntu Advantage Cloud Guest:
http://www.ubuntu.com/business/services/cloud

13 packages can be updated.
8 updates are security updates.
```

System restart required Last login: Mon Apr 4 18:58:15 2016 from 71-222-65-56.ptld.qwest.net don@cloud:~\$

To update the packages use apt-get update to refresh the package lists. You will need to escalate privileges (become root)

```
don@cloud:~$ sudo bash
[sudo] password for don:
Sorry, try again.
[sudo] password for don:
root@cloud:~# apt-get update
Ign http://us-west-2.ec2.archive.ubuntu.com trusty InRelease
Get:1 http://us-west-2.ec2.archive.ubuntu.com trusty-updates InRelease [65.9 kB]
Hit http://us-west-2.ec2.archive.ubuntu.com trusty Release.gpg
Hit http://us-west-2.ec2.archive.ubuntu.com trusty Release
Get:2 http://us-west-2.ec2.archive.ubuntu.com trusty-updates/main Sources [271 kB]
Get:3 http://us-west-2.ec2.archive.ubuntu.com trusty-updates/universe Sources [152 kB]
Get:4 http://us-west-2.ec2.archive.ubuntu.com trusty-updates/main amd64 Packages [753 kB]
Get:5 http://us-west-2.ec2.archive.ubuntu.com trusty-updates/universe amd64 Packages [358 kB]
Hit http://us-west-2.ec2.archive.ubuntu.com trusty-updates/main Translation-en
Hit http://us-west-2.ec2.archive.ubuntu.com trusty-updates/universe Translation-en
Hit http://us-west-2.ec2.archive.ubuntu.com trusty/main Sources
Hit http://us-west-2.ec2.archive.ubuntu.com trusty/universe Sources
Hit http://us-west-2.ec2.archive.ubuntu.com trusty/main amd64 Packages
Hit http://us-west-2.ec2.archive.ubuntu.com trusty/universe amd64 Packages
Hit http://us-west-2.ec2.archive.ubuntu.com trusty/main Translation-en
Hit http://us-west-2.ec2.archive.ubuntu.com trusty/universe Translation-en
Ign http://us-west-2.ec2.archive.ubuntu.com trusty/main Translation-en_US
Ign http://us-west-2.ec2.archive.ubuntu.com trusty/universe Translation-en_US
Get:6 http://security.ubuntu.com trusty-security InRelease [65.9 kB]
Get:7 http://security.ubuntu.com trusty-security/main Sources [110 kB]
Get:8 http://security.ubuntu.com trusty-security/universe Sources [35.2 kB]
Get:9 http://security.ubuntu.com trusty-security/main amd64 Packages [455 kB]
Get:10 http://security.ubuntu.com trusty-security/universe amd64 Packages [126 kB]
Hit http://security.ubuntu.com trusty-security/main Translation-en
Hit http://security.ubuntu.com trusty-security/universe Translation-en
Fetched 2,393 kB in 3s (679 kB/s)
Reading package lists... Done
root@cloud:~#
```

once this is done you can update all of the installed packages to the currently supported versions by using the dist-upgrade command.

```
root@cloud:~# apt-get dist-upgrade
Reading package lists... Done
Building dependency tree
Reading state information... Done
Calculating upgrade... Done
The following packages were automatically installed and are no longer required:
  linux-headers-3.13.0-76 linux-headers-3.13.0-76-generic
  linux-headers-3.13.0-77 linux-headers-3.13.0-77-generic
  linux-image-3.13.0-76-generic linux-image-3.13.0-77-generic
Use 'apt-get autoremove' to remove them.
The following NEW packages will be installed:
  linux-headers-3.13.0-85 linux-headers-3.13.0-85-generic
  linux-image-3.13.0-85-generic
The following packages will be upgraded:
  apt apt-transport-https apt-utils libapt-inst1.5 libapt-pkg4.12 libpq5
  linux-headers-generic linux-headers-virtual linux-image-virtual
  linux-libc-dev linux-virtual postgresql-9.3 postgresql-client-9.3
  postgresql-contrib-9.3 postgresql-doc-9.3
15 upgraded, 3 newly installed, 0 to remove and 0 not upgraded.
```

```

Need to get 33.2 MB of archives.
After this operation, 120 MB of additional disk space will be used.
Do you want to continue? [Y/n]
Get:1 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main libapt-pkg4.12 amd64 1.0.1ubuntu2.12 [637 kB]
Get:2 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main apt amd64 1.0.1ubuntu2.12 [954 kB]
Get:3 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main libapt-inst1.5 amd64 1.0.1ubuntu2.12 [58.6 kB]
Get:4 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-image-3.13.0-85-generic amd64 3.13.0-85.129 [15.2 MB]
Get:5 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main apt-utils amd64 1.0.1ubuntu2.12 [172 kB]
Get:6 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main apt-transport-https amd64 1.0.1ubuntu2.12 [25.1 kB]
Get:7 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main libpq5 amd64 9.3.12-0ubuntu0.14.04 [78.5 kB]
Get:8 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-headers-3.13.0-85-all 3.13.0-85.129 [8,887 kB]
Get:9 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-headers-3.13.0-85-generic amd64 3.13.0-85.129 [707 kB]
Get:10 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-virtual amd64 3.13.0-85.91 [1,778 B]
Get:11 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-image-virtual amd64 3.13.0-85.91 [2,240 B]
Get:12 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-headers-virtual amd64 3.13.0-85.91 [1,756 B]
Get:13 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-headers-generic amd64 3.13.0-85.91 [2,230 B]
Get:14 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main linux-libc-dev amd64 3.13.0-85.129 [775 kB]
Get:15 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main postgresql-contrib-9.3 amd64 9.3.12-0ubuntu0.14.04 [401 kB]
Get:16 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main postgresql-client-9.3 amd64 9.3.12-0ubuntu0.14.04 [785 kB]
Get:17 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main postgresql-9.3 amd64 9.3.12-0ubuntu0.14.04 [2,691 kB]
Get:18 http://us-west-2.ec2.archive.ubuntu.com/ubuntu/ trusty-updates/main postgresql-doc-9.3-all 9.3.12-0ubuntu0.14.04 [1,780 kB]
Fetched 33.2 MB in 0s (36.5 MB/s)
(Reading database ... 163885 files and directories currently installed.)
Preparing to unpack .../libapt-pkg4.12_1.0.1ubuntu2.12_amd64.deb ...
Unpacking libapt-pkg4.12:amd64 (1.0.1ubuntu2.12) over (1.0.1ubuntu2.11) ...
Setting up libapt-pkg4.12:amd64 (1.0.1ubuntu2.12) ...
Processing triggers for libc-bin (2.19-0ubuntu6.7) ...
(Reading database ... 163885 files and directories currently installed.)
Preparing to unpack .../apt_1.0.1ubuntu2.12_amd64.deb ...
Unpacking apt (1.0.1ubuntu2.12) over (1.0.1ubuntu2.11) ...
Processing triggers for man-db (2.6.7.1-1ubuntu1) ...
Setting up apt (1.0.1ubuntu2.12) ...
Processing triggers for libc-bin (2.19-0ubuntu6.7) ...
(Reading database ... 163885 files and directories currently installed.)
Preparing to unpack .../libapt-inst1.5_1.0.1ubuntu2.12_amd64.deb ...
Unpacking libapt-inst1.5:amd64 (1.0.1ubuntu2.12) over (1.0.1ubuntu2.11) ...
Selecting previously unselected package linux-image-3.13.0-85-generic.
Preparing to unpack .../linux-image-3.13.0-85-generic_3.13.0-85.129_amd64.deb ...
Done.
Unpacking linux-image-3.13.0-85-generic (3.13.0-85.129) ...
Preparing to unpack .../apt-utils_1.0.1ubuntu2.12_amd64.deb ...
Unpacking apt-utils (1.0.1ubuntu2.12) over (1.0.1ubuntu2.11) ...
Preparing to unpack .../apt-transport-https_1.0.1ubuntu2.12_amd64.deb ...
Unpacking apt-transport-https (1.0.1ubuntu2.12) over (1.0.1ubuntu2.11) ...
Preparing to unpack .../libpq5_9.3.12-0ubuntu0.14.04_amd64.deb ...
Unpacking libpq5 (9.3.12-0ubuntu0.14.04) over (9.3.11-0ubuntu0.14.04) ...
Selecting previously unselected package linux-headers-3.13.0-85.
Preparing to unpack .../linux-headers-3.13.0-85_3.13.0-85.129_all.deb ...
Unpacking linux-headers-3.13.0-85 (3.13.0-85.129) ...
Selecting previously unselected package linux-headers-3.13.0-85-generic.
Preparing to unpack .../linux-headers-3.13.0-85-generic_3.13.0-85.129_amd64.deb ...
Unpacking linux-headers-3.13.0-85-generic (3.13.0-85.129) ...
Preparing to unpack .../linux-virtual_3.13.0-85.91_amd64.deb ...
Unpacking linux-virtual (3.13.0-85.91) over (3.13.0-83.89) ...
Preparing to unpack .../linux-image-virtual_3.13.0-85.91_amd64.deb ...
Unpacking linux-image-virtual (3.13.0-85.91) over (3.13.0-83.89) ...
Preparing to unpack .../linux-headers-virtual_3.13.0-85.91_amd64.deb ...
Unpacking linux-headers-virtual (3.13.0-85.91) over (3.13.0-83.89) ...
Preparing to unpack .../linux-headers-generic_3.13.0-85.91_amd64.deb ...
Unpacking linux-headers-generic (3.13.0-85.91) over (3.13.0-83.89) ...
Preparing to unpack .../linux-libc-dev_3.13.0-85.129_amd64.deb ...
Unpacking linux-libc-dev:amd64 (3.13.0-85.129) over (3.13.0-83.127) ...
Preparing to unpack .../postgresql-contrib-9.3_9.3.12-0ubuntu0.14.04_amd64.deb ...
Unpacking postgresql-contrib-9.3 (9.3.12-0ubuntu0.14.04) over (9.3.11-0ubuntu0.14.04) ...
Preparing to unpack .../postgresql-client-9.3_9.3.12-0ubuntu0.14.04_amd64.deb ...
Unpacking postgresql-client-9.3 (9.3.12-0ubuntu0.14.04) over (9.3.11-0ubuntu0.14.04) ...
Preparing to unpack .../postgresql-9.3_9.3.12-0ubuntu0.14.04_amd64.deb ...
* Stopping PostgreSQL 9.3 database server
Unpacking postgresql-9.3 (9.3.12-0ubuntu0.14.04) over (9.3.11-0ubuntu0.14.04) ...
Preparing to unpack .../postgresql-doc-9.3_9.3.12-0ubuntu0.14.04_all.deb ...
Unpacking postgresql-doc-9.3 (9.3.12-0ubuntu0.14.04) over (9.3.11-0ubuntu0.14.04) ...
Processing triggers for man-db (2.6.7.1-1ubuntu1) ...
Processing triggers for postgresql-common (154ubuntu1) ...
Building PostgreSQL dictionaries from installed myspell/hunspell packages...
Removing obsolete dictionary files:
Setting up libapt-inst1.5:amd64 (1.0.1ubuntu2.12) ...
Setting up linux-image-3.13.0-85-generic (3.13.0-85.129) ...
Running depmod.
update-initramfs: deferring update (hook will be called later)
Examining /etc/kernel/postinst.d.
run-parts: executing /etc/kernel/postinst.d/apt-auto-removal 3.13.0-85-generic /boot/vmlinuz-3.13.0-85-generic
run-parts: executing /etc/kernel/postinst.d/initramfs-tools 3.13.0-85-generic /boot/vmlinuz-3.13.0-85-generic
update-initramfs: Generating /boot/initrd.img-3.13.0-85-generic
run-parts: executing /etc/kernel/postinst.d/update-notifier 3.13.0-85-generic /boot/vmlinuz-3.13.0-85-generic
run-parts: executing /etc/kernel/postinst.d/x-grub-legacy-ec2 3.13.0-85-generic /boot/vmlinuz-3.13.0-85-generic
Searching for GRUB installation directory ... found: /boot/grub
Searching for default file ... found: /boot/grub/default
Testing for an existing GRUB menu.lst file ... found: /boot/grub/menu.lst
Searching for splash image ... none found, skipping ...
Found kernel: /boot/vmlinuz-3.13.0-85-generic
Found kernel: /boot/vmlinuz-3.13.0-83-generic
Found kernel: /boot/vmlinuz-3.13.0-79-generic
Found kernel: /boot/vmlinuz-3.13.0-77-generic
Found kernel: /boot/vmlinuz-3.13.0-76-generic

```

[OK]


```

Found kernel: /boot/vmlinuz-3.13.0-74-generic
Replacing config file /run/grub/menu.lst with new version
Updating /boot/grub/menu.lst ... done

run-parts: executing /etc/kernel/postinst.d/zz-update-grub 3.13.0-85-generic /boot/vmlinuz-3.13.0-85-generic
Generating grub configuration file ...
Found linux image: /boot/vmlinuz-3.13.0-85-generic
Found initrd image: /boot/initrd.img-3.13.0-85-generic
Found linux image: /boot/vmlinuz-3.13.0-83-generic
Found initrd image: /boot/initrd.img-3.13.0-83-generic
Found linux image: /boot/vmlinuz-3.13.0-79-generic
Found initrd image: /boot/initrd.img-3.13.0-79-generic
Found linux image: /boot/vmlinuz-3.13.0-77-generic
Found initrd image: /boot/initrd.img-3.13.0-77-generic
Found linux image: /boot/vmlinuz-3.13.0-76-generic
Found initrd image: /boot/initrd.img-3.13.0-76-generic
Found linux image: /boot/vmlinuz-3.13.0-74-generic
Found initrd image: /boot/initrd.img-3.13.0-74-generic
done
Setting up apt-utils (1.0.1ubuntu2.12) ...
Setting up apt-transport-https (1.0.1ubuntu2.12) ...
Setting up libpq5 (9.3.12-0ubuntu0.14.04) ...
Setting up linux-headers-3.13.0-85 (3.13.0-85.129) ...
Setting up linux-headers-3.13.0-85-generic (3.13.0-85.129) ...
Setting up linux-image-virtual (3.13.0.85.91) ...
Setting up linux-headers-generic (3.13.0.85.91) ...
Setting up linux-headers-virtual (3.13.0.85.91) ...
Setting up linux-virtual (3.13.0.85.91) ...
Setting up linux-libc-dev:amd64 (3.13.0-85.129) ...
Setting up postgresql-client-9.3 (9.3.12-0ubuntu0.14.04) ...
Setting up postgresql-9.3 (9.3.12-0ubuntu0.14.04) ...
* Starting PostgreSQL 9.3 database server
Setting up postgresql-contrib-9.3 (9.3.12-0ubuntu0.14.04) ...
Setting up postgresql-doc-9.3 (9.3.12-0ubuntu0.14.04) ...
Processing triggers for libc-bin (2.19-0ubuntu6.7) ...

```

[OK]

If one of the updates includes a kernel or if system restart required a system reboot should be scheduled.

5.88 TaskAddGitHubRepo

5.88.1 Add "Trac"king (and Mirroring) to Github Repos

Trac supports git repos on the local machine which can be extended to GitHub repos by using the provided post commit hook. Adding the repos is a little convoluted but once set up all commits to the GitHub repository cause a backup copy to be made on the local server.

There are two components at work here. The built in git/svn repository browser and the GitHub plugin. The GitHub plugin provides the webhook.

Process

In this example we are going to add one of suspect devices repos to the git server and add a webhook to synchronize the two.

[[Image(TaskAddGitHubRepo:github-example.png,100%)]] Ssh in and clone the repo to the git/trac server. (repos are at /var/trac/devel/repos/)

To make life less of a pain www-data is set up so that you can su to the account.

```
haifisch:~ don$ ssh feurig@git
...
feurig@douglas:~$ sudo bash
[sudo] password for feurig:
root@douglas:~# su - www-data
www-data@douglas:~$ cd /var/trac/devel/repos/
www-data@douglas:/var/trac/devel/repos$ git clone --mirror git@github.com:suspect-devices/errata_physical_computing.git
Cloning into bare repository 'errata_physical_computing.git'...
remote: Enumerating objects: 26, done.
remote: Total 26 (delta 0), reused 0 (delta 0), pack-reused 26
Receiving objects: 100% (26/26), 19.48 KiB | 19.48 MiB/s, done.
Resolving deltas: 100% (8/8), done.
www-data@douglas:/var/trac/devel/repos$
```

Log into the [git/trac site](#) and navigate to admin->repositories [[Image(TaskAddGitHubRepo:add-repo.png,100%)]] Add the repo. [[Image(TaskAddGitHubRepo:repo-added.png,100%)]] Copy the command presented and execute it as www-data.

```
www-data@douglas:/var/trac/devel/repos$ trac-admin "/var/trac/devel/env" repository resync "PC-errata"
Resyncing repository history for PC-errata...
0 revisions cached.
PC-errata is not a cached repository.
Done.
www-data@douglas:/var/trac/devel/repos$
```

Check the the GitHub web hook url by browsing git.suspectdevices.com/devel/github/

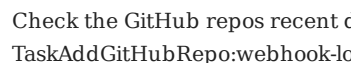

It should return the following: "Endpoint is ready to accept GitHub notifications." [[Image(TaskAddGitHubRepo:endpoint-ready.png,100%)]]

Enable the web hook on your GitHub repository. *Navigate to the settings -> webhooks -> Add Webhook* * Paste the url you just tested * Make sure that you select json * Disable checking the certificate since its self signed. *(it will bitch)* [[Image(TaskAddGitHubRepo:add-webhook.png,100%)]]

Make changes to the repository.

```
haifisch:~ don$ cd /tmp/
haifisch:tmp don$ git clone git@github.com:suspect-devices/errata_physical_computing.git
Cloning into 'errata_physical_computing'...
remote: Enumerating objects: 26, done.
remote: Total 26 (delta 0), reused 0 (delta 0), pack-reused 26
Receiving objects: 100% (26/26), 19.48 KiB | 6.49 MiB/s, done.
Resolving deltas: 100% (8/8), done.
haifisch:tmp don$ cd errata_physical_computing/
haifisch:errata_physical_computing don$ nano proof.md
haifisch:errata_physical_computing don$ git commit -a -m "Fix weird cruft at the top of the markdown"
[master 2ccl077] Fix weird cruft at the top of the markdown
1 file changed, 1 insertion(+), 1 deletion(-)
haifisch:errata_physical_computing don$ git push
Counting objects: 3, done.
Delta compression using up to 4 threads.
Compressing objects: 100% (3/3), done.
Writing objects: 100% (3/3), 318 bytes | 318.00 KiB/s, done.
```

```
Total 3 (delta 2), reused 0 (delta 0)
remote: Resolving deltas: 100% (2/2), completed with 2 local objects.
To github.com:suspect-devices/errata_physical_computing.git
 57401a6..2cc1d07  master -> master
```

Check the GitHub repos recent deliveries at the bottom of Settings->Webhooks->Manage Webhook  Browse the code changes on the git/trac server. 

5.88.2 References

- <https://serverdocs.suspectdevices.com/tracdocs/wiki/TracRepositoryAdmin>
- <https://github.com/trac-hacks/trac-github>

5.89 TaskAddLxdContainerWithAnsible

5.89.1 New Container Using Ansible

With Ansible added to kb2018 we expand on the profiles we use to create users and create a sane environment. There are two steps required to create a container on kb2018.

1. Add the name, ip_address, and purpose to the inventory file `/etc/ansible/hosts`.

```
... redshirt ip_address=198.202.31.200 purpose="Disposable Ubuntu" ...
```

2. Run the ansible playbook `/etc/ansible/playbooks/create-lxd-containers.yml`

```
root@kb2018:/etc/ansible# ansible-playbook /etc/ansible/playbooks/create-lxd-containers.yml
```

```
PLAY [localhost] *****
```

If you want something other than ubuntu-lts you can: * set the image_alias. *these are images that we know work in our environment*

```
root@kb2018:/etc/ansible# lxc image list
+-----+-----+-----+-----+-----+-----+-----+-----+
| ALIAS | FINGERPRINT | PUBLIC | DESCRIPTION | ARCH | SIZE | UPLOAD DATE |
+-----+-----+-----+-----+-----+-----+-----+
| centos/7c | 700c86f31546 | no | Centos 7 (20190109_02:16) plus cloud | x86_64 | 172.49MB | Mar 21, 2019 at 1:54am (UTC) |
+-----+-----+-----+-----+-----+-----+-----+
| debian/9c | 38d17964647d | no | Debian stretch (20190108_05:24) plus cloud | x86_64 | 227.58MB | Mar 19, 2019 at 5:57am (UTC) |
+-----+-----+-----+-----+-----+-----+-----+
| ubuntu-lts | c395a7105278 | no | ubuntu 18.04 LTS amd64 (release) (20180911) | x86_64 | 173.98MB | Sep 29, 2018 at 11:50pm (UTC) |
+-----+-----+-----+-----+-----+-----+-----+
root@kb2018:~# cd /etc/ansible/
root@kb2018:/etc/ansible# ls
ansible.cfg files hosts host_vars playbooks README.md roles
oot@kb2018:/etc/ansible# lxc delete redshirt --force
root@kb2018:/etc/ansible# nano hosts
...
redshirt ip_address=198.202.31.200 purpose="Disposable Debian" image_alias="debian/9c"
...
```

And (re)run the playbook.

```
root@kb2018:/etc/ansible# ansible-playbook /etc/ansible/playbooks/create-lxd-containers.yml
```

You can also create infrastructure servers by setting `net_and_disk_profile` to "infra".

The ansible playbook and host file are maintained in a private bitbucket repository. If you add roles or create a host that you want to keep please update the repository. *Ignore the errors, I will reconfigure a user for kb2018 when bitbucket really stops supporting the organization account*

```
feurig@kb2018:~$ sudo bash
[sudo] password for feurig:
root@kb2018:~# cd /etc/ansible/hosts
root@kb2018:/etc/ansible# nano hosts
...
morgan ip_address=198.202.31.224 purpose="Infrastructure Test Machine" net_and_disk_profile="infra"
...
root@kb2018:/etc/ansible# ansible-playbook /etc/ansible/playbooks/create-lxd-containers.yml
PLAY RECAP
*****
localhost : ok=4 changed=1 unreachable=0 failed=0

root@kb2018:/etc/ansible# git commit -a -m "Recreate Morgan as Infrastructure Test Server"
[master 10d4ce0] Recreate Morgan as Infrastructure Test Server
Committer: Root at KB2018 <root@kb2018.suspectdevices.com>
...
1 file changed, 1 insertion(+), 1 deletion(-)
root@kb2018:/etc/ansible# git push
Counting objects: 3, done.
Delta compression using up to 16 threads.
Compressing objects: 100% (3/3), done.
Writing objects: 100% (3/3), 378 bytes | 378.00 KiB/s, done.
Total 3 (delta 2), reused 0 (delta 0)
remote:
remote: Warning!
remote: You are currently connecting with your team account.
remote: This is no longer supported, so please connect using your user account.
remote:
To bitbucket.org:suspectdevicesadmin/ansible.git
```

```
d7f5f12..10d4ce0 master -> master  
root@kb2018:/etc/ansible#
```

5.90 TaskCreatingNewContainers

5.90.1 Creating containers

LXD allows us to create lightweight virtual machines, and combined with filesystems such as ZFS, provides several mechanisms to easily configure, backup, replicate and update them. Adding ansible to the mix makes this process even more simple.

To create a container in our current environment simply add the hostname, ip_address, and purpose to /etc/ansible/hosts and run the create-lxd-containers.yml

```
root@kb2018:/etc/ansible# nano hosts
...
redshirt ip_address=198.202.31.200 purpose="Disposable Ubuntu"
...
root@kb2018:/etc/ansible# ansible-playbook /etc/ansible/playbooks/create-lxd-containers.yml
```

The containers created have admin accounts and ssh keys installed. They have Isolated static ip addresses which can not reach the server directly. Getting onto the containers requires an ssh key and a password to escalate privileges. By default the containers are unprivileged which should minimize the security risks to the main server. They also have an os agnostic script to perform periodic updates.

Mechanisms

LXD provides images and profiles which define the disk storage, network and other configuration used to create the container. The profiles include cloud config however only the ubuntu image implements it. Most things work well except some don't or some things change and the updates or the updated images require some tweaking (*resolved for instance because it wasn't broken and you had to work around what was already*)

5.90.2 Container !Image/Profile notes

The fragments below are from my work to create other images that would work as well.

Ubuntu LTS (20.04)

My initial forays into the new LTS release are [StartUsingTheFwords here].

Previous Ubuntu LTS (18.04)

Ubuntu 18.04 is really well suited for LXD in that it comes stock with cloud init. This means that with a simple profile you can create a usable container pre seeded with admin accounts, static networking and an update script.

Creating a new ubuntu container using lxd

In our environment hosts do not tend to be temporary so they are not dynamically allocated. Containers are built at the static ip address at redshirt.suspectdevices.com and then reconfigured before being used.

```
root@bs2020:~# lxc image list kb2018:
+-----+-----+-----+-----+-----+-----+-----+-----+
| ALIAS | FINGERPRINT | PUBLIC | DESCRIPTION | ARCH | SIZE | UPLOAD DATE |
+-----+-----+-----+-----+-----+-----+-----+
| ubuntu-lts | ae465acff89b | no | ubuntu 18.04 LTS amd64 (release) (20180613) | x86_64 | 173.14MB | Jun 16, 2018 at 10:07pm (UTC) |
+-----+-----+-----+-----+-----+-----+-----+
root@bs2020:~# lxc init kb2018:ubuntu-lts test18 -p susdev19 -p default
Creating test18
root@bs2020:~# lxc start test18
root@bs2020:~# lxc exec test18 bash
root@test18:~# nano /etc/netplan/50-cloud-init.yaml
network:
  version: 2
  ethernets:
    eth0:
      dhcp4: no
      addresses: [198.202.31.216/25]
      gateway4: 198.202.31.129
      nameservers:
```

```

search: [suspectdevices.com fromhell.com vpn]
addresses: [198.202.31.141]
root@test18:~# netplan apply
root@test18:~# update.sh
root@test18:~# reboot
root@test18:~# root@bs2020:~# lxc list
+-----+-----+-----+-----+-----+-----+
| NAME | STATE | IPV4 | IPV6 | TYPE | SNAPSHOTS |
+-----+-----+-----+-----+-----+-----+
....
+-----+-----+-----+-----+-----+-----+
| test18 | RUNNING | 198.202.31.216 (eth0) | | PERSISTENT | 0 |
+-----+-----+-----+-----+-----+-----+
root@bs2020:~#

```

lxd profile for suspectdevices

The profile for suspect devices is broken into three major parts. * Network configuration * System, User and Security Configuration * Disk Pools and Network Devices

OVERRIDING THE PROFILES NETWORK CONFIGURATION

.... do this by hand and discuss the ansible update in progress....

```

version: 1
config:
- type: physical
  name: eth0
  subnets:
  - type: static
    ipv4: true
    address: 198.202.31.200
    netmask: 255.255.255.128
    gateway: 198.202.31.129
    control: auto
- type: nameserver
  address: 198.202.31.141

```

System, User and Security Configuration

(Network Configuration is Stubbed in)

```

root@kb2018:~# lxc profile show susdev19
config:
  user.network-config: |
    version: 1
    config:
      - type: physical
        name: eth0
        subnets:
          - type: static
            ipv4: true
            address: 198.202.31.200
            netmask: 255.255.255.128
            gateway: 198.202.31.129
            control: auto
      - type: nameserver
        address: 198.202.31.141
  user.user-data: |
    #cloud-config
    timezone: America/Vancouver
    users:
      - name: feurig
        passwd: "$6$2Pf0ittl$nl....VdI/FyCXtu."
        geccos: Donald Delmar Davis
        ssh-authorized-keys:
          - ssh-rsa AAAA....uj4SL don@annie
          - ssh-rsa AAA.....FMNNn don@haifisch.local
        groups: sudo,root,wheel
        shell: /bin/bash
      - name: joe
        passwd: "$6$o14Dp3u...pD3vLrSlvX."
        geccos: Joseph Wayne Dumoulin
        ssh-authorized-keys:
          - ssh-rsa AAAA...r6Y/ZepPr jdumoulin@nextit.com
        groups: sudo,root,wheel
        shell: /bin/bash
  manage_resolv_conf: false
  packages:
  - python
  package_update: true
  package_upgrade: true
  write_files:
  - path: /etc/systemd/resolved.conf
    permissions: '0644'

```

```

owner: root:root
content: |
  # resolved because that wasnt broken either
  [Resolve]
  DNS= 198.202.31.141 198.202.31.132 8.8.4.4
- path: /usr/local/bin/update.sh
permissions: '0774'
owner: root:root
content: |
  #!/bin/bash
  # update.sh for debian/ubuntu/centos (copyleft) don@suspecdevices.com
  echo ----- begin updating `uname -n` -----
  if [ -x "$(command -v apt-get)" ]; then
    apt-get update
    apt-get -y dist-upgrade
    apt-get -y autoremove
  fi
  if [ -x "$(command -v yum)" ]; then
    echo yum upgrade.
    yum -y upgrade
  fi
  if [ -x "$(command -v zypper)" ]; then
    echo zypper dist-upgrade.
    zypper -y dist-upgrade
  fi
  echo =====### done=====
runcmd:
  # fix stupid subtle things
  # sudo needs to be able to resolve itself to authenticate users
  # and the users are locked by default
  # cloud cart blanch accounts are inexcusable
  - sed -i "s/^127.0.0.1/#127.0.0.1/" /etc/hosts
  - echo 127.0.0.1 `hostname` localhost >>/etc/hosts
  - passwd joe -u
  - passwd feurig -u
  - userdel -f ubuntu
  - userdel -f centos
  - userdel -f opensuse
  #- netplan apply
power_state:
  mode: reboot
  message: See You Soon...
  condition: True
description: Try to create a sane environment for cloud-init based operating systems
devices: {}
name: susdev19
used_by:

```

5.90.3 cloud - init and ubuntu but no where else.

When I ran up the lxc containers for operating systems that aren't ubuntu the magic profile doesn't work. And in fact cloud-init and the utilities that are native on Ubuntu are not installed on those images. (WTF??) So my first attempt (using the Debian 9 container was to add cloud-init cloud-utils and the other packages and then get export and reimport the container) which more or less failed miserably (because creating new containers from scratch isn't as simple as they say it is. :).

debian 9 (Works!)

It turns out I didn't need to export or import the image. I just needed to copy the lxc templates from a working ubuntu image and then modify metadata.yaml on the image while its running and publish the result. (this method is buried in the discussion [| here](#))

```

... lxc create using images:debian/9 ....
... lxc start image and add cloud init and cloud utils ...
... copy templates and metadata data from working ubuntu ....
... link /etc/network/interfaces.d/50... -> /etc/network/interfaces ...
... delete /var/log/cloud cruft ...
... shutdown and lxc publish ....
root@bs2020:~# lxc publish kernigan --alias debian/9c
root@bs2020:~# lxc init debian/9c redshirt -p susdev19

```

This works well. More better documentation to follow.

Centos 7 (works)

```

[root@keynes ~]# cd /etc/sysconfig/
[root@keynes sysconfig]# cat network
NETWORKING=yes
HOSTNAME=LXC_NAME
[root@keynes sysconfig]# cd network-scripts/
[root@keynes network-scripts]# vi ifcfg-eth0
DEVICE=eth0
BOOTPROTO=none

```



```

ONBOOT=yes
HOSTNAME=LXC_NAME
NM_CONTROLLED=no
TYPE=Ethernet
PREFIX=25
IPADDR=198.202.31.220
MTU=
GATEWAY=198.202.31.129
[root@keynes network-scripts]# systemctl restart network
[root@keynes network-scripts]# nano /etc/resolv.conf
bash: nano: command not found
[root@keynes network-scripts]# vi /etc/resolv.conf
...
nameserver 198.202.31.141
search suspectdevices.com
...
[root@keynes network-scripts]# ping digithink.com
PING digithink.com (198.202.31.230) 56(84) bytes of data.
64 bytes from 198.202.31.230 (198.202.31.230): icmp_seq=1 ttl=64 time=0.441 ms
...
[root@keynes network-scripts]# cd
[root@keynes ~]# yum update
Failed to set locale, defaulting to C
Loaded plugins: fastestmirror
...
updates | 3.4 kB
00:00:00
(1/4): extras/7/x86_64/primary_db | 156 kB
00:00:00
(2/4): updates/7/x86_64/primary_db | 1.3 MB
00:00:00
(3/4): base/7/x86_64/group_gz | 166 kB
00:00:00
(4/4): base/7/x86_64/primary_db | 6.0 MB
00:00:01
No packages marked for update
[root@keynes ~]# yum install -y nano less
[root@keynes ~]# yum install -y cloud-init
[root@keynes ~]# yum install -y cloud-utils
[root@keynes ~]# yum install -y openssh-server
[root@keynes ~]# yum install -y sudo
[root@keynes ~]# cat >>/etc/sudoers.d/9_fix-centos-sudo <<EOD
%sudo ALL=(ALL) ALL
centos ALL = /usr/bin/su nobody
EOD
[root@keynes ~]# exit
... modify metadata.yaml ...
... copy templates ....
root@bs2020:~# lxc image delete centos/7c
root@bs2020:~# lxc publish keynes --alias centos/7c
Container published with fingerprint: a27609a23021f4577dfea987176fa942635d349b2e3be0e046118db88af4c56a
root@bs2020:~#
root@bs2020:~# lxc launch centos/7c redshirt -p susdev19
Creating redshirt
Starting redshirt

```

Check the work....

```

haifisch:~ don$ ssh feurig@redshirt.suspectdevices.com
The authenticity of host 'redshirt.suspectdevices.com (198.202.31.200)' can't be established.
ECDSA key fingerprint is SHA256:ad0/DY7qDL9XKS14lnjSq9jv63e18Nrr4IZjT0yu70g.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'redshirt.suspectdevices.com,198.202.31.200' (ECDSA) to the list of known hosts.
[feurig@redshirt ~]$ sudo bash

We trust you have received the usual lecture from the local System
Administrator. It usually boils down to these three things:

    #1) Respect the privacy of others.
    #2) Think before you type.
    #3) With great power comes great responsibility.

[sudo] password for feurig:
[root@redshirt feurig]#

```

todo: look deeper into right way to sudo and cloud config. (ubuntu:ubuntu problem)

Fedora 29 (not picking up local cloud-init)

```

[root@kundara ~]# cat .bash_history
ip address add 198.202.31.200/25 dev eth0
ip route add default via 198.202.31.129
ping digithink.com
dnf upgrade
dnf upgrade cloud-init
dnf install cloud-init
dnf install -y cloud-utils
dnf install -y nano less sudo
dnf install -y openssh-server

```

```
cat >>/etc/sudoers.d/9_fix-fedora-sudo <<EOD
%sudo ALL=(ALL) ALL
fedora ALL = /usr/bin/su nobody
EOD
dnf install -y network-scripts
echo "NOZEROCONF=yes" >> /etc/sysconfig/network
systemctl enable cloud-init
chkconfig --levels 2345 sshd on
chkconfig --levels 2345 network on
journalctl --vacuum-time='date +%s'
shutdown -h now
```

OpenSuse 15.0 (Close -- some oddities)

Image comes up with no network.

```
ip address add 198.202.31.200/25 dev eth0
ip route add default via 198.202.31.129
cat >/etc/resolv.conf<<EOD
nameserver 198.202.31.141
nameserver 198.202.31.132
nameserver 8.8.8.8
search suspectdevices.com
EOD
```

Install the packages needed to work here.

```
zypper -y install nano sudo cloud-init
zypper -y install openssh
zypper -y dist-upgrade

systemctl enable cloud-init
```

sshd install masks the service as disabled.

```
systemctl unmask sshd
systemctl enable sshd
```

Sudo comes out of the box pretty insecurely configured.

```
cat> /etc/sudoers.d/9_fix_opensuse_sudo<<EOD
Defaults !targetpw
%sudo ALL=(ALL) ALL
opensuse ALL = /usr/bin/su nobody
EOD
```

todo: (cloud init bugs) * figure out why hashed passwords don't work. (or is it just my long complicated password). * figure out why the default route isn't getting propagated.

updating running containers

The update script created by the profile can be easily executed on all running containers on both hosts with the following 2 lines of bash.

```
root@kb2018:~# for h in `lxc list bs2020: -c n --format csv` ;do echo $h ;lxc exec bs2020:$h update.sh; done
root@kb2018:~# for h in `lxc list local: -c n --format csv` ;do echo $h ;lxc exec local:$h update.sh; done
```

5.91 Task: Dual Proxy Configuration

... you are here ...

5.91.1 Link Dump

- <https://www.christianschenk.org/blog/using-a-parent-proxy-with-squid/>
- <https://stackoverflow.com/questions/21886716/lightweight-forwarding-proxy-with-auth-support>
- <https://itandsecuritystuffs.wordpress.com/2015/01/22/how-install-a-proxy-server-to-anonymise-your-internet-surfing/>
- <https://www.privoxy.org/user-manual/index.html>
- <https://www.neowin.net/forum/topic/601824-need-a-http-proxy-server-that-supports-socks-parent/>
- <https://sourceforge.net/p/ijbswa/mailman/message/19931928/>
- <https://github.com/crozuk/pi-hole-wireguard-privoxy>

5.92 Fast Forward

Installing Debian packages from newer/previous distributions One of the compromises made in Ubuntu's long term support release cycle is that stability is preferred over features. This is usually a good thing however sometimes you need features that are only found in a future release. Two examples of this are trac-1.2.2 which has a working git integration, which is broken in 18.04's version (trac-1.2). Another is GNUCobol's "Stable" version (2.2).

5.92.1 Manual installation

For trac, I pulled the package file from 18.10's repositories and installed it manually. This breaks any updates or security fixes that are made to the newer repository, as well as the base ones. I don't much care for this solution and won't map it out here.

5.92.2 Adding Future Repositories.

Adding future repositories to /etc/apt/sources allows us to pull from those repositories.

```
root@redshirt:~# nano /etc/apt/sources.list
....
deb http://archive.ubuntu.com/ubuntu/ disco restricted main multiverse universe
deb http://archive.ubuntu.com/ubuntu/ disco-updates restricted main multiverse universe
deb http://security.ubuntu.com/ubuntu/ disco-security restricted main multiverse universe
```

Unfortunately the newer repo now becomes the default repo for everything in it. Essentially, the next apt-get dist-upgrade will take your entire install to the bleeding edge.

5.92.3 Google sucks

(AKA Following the instructions on <https://medium.com/@george.shuklin/how-to-install-packages-from-a-newer-distribution-without-installing-unwanted-6584fa93208f>) In addition to trying to get you to give your facebook or google credentials the top listed instructions on installing specific packages don't work. It did, however, provide clues.

More or less they tell you create the following file (somefile.pref) in /etc/apt/preferences.d/ and run apt-get update.

```
Package: *
Pin: release n=disco
Pin-Priority: -10
Package: gnucobol
Pin: release n=disco
Pin-Priority: 500
```

Following these instructions caused disco-security and disco-updates to have the same priority as the current release (bionic).

```
root@redshirt:~# nano /etc/apt/preferences.d/gnucobol22.pref
...
root@redshirt:~# apt-get update
...
root@redshirt:~# apt-cache policy
Package files:
 100 /var/lib/dpkg/status
    release a=now
 500 http://security.ubuntu.com/ubuntu disco-security/universe amd64 Packages
    release v=19.04,o=Ubuntu,a=disco-security,n=disco,l=Ubuntu,c=universe,b=amd64
...

 500 http://archive.ubuntu.com/ubuntu bionic/main amd64 Packages
    release v=18.04,o=Ubuntu,a=bionic,n=bionic,l=Ubuntu,c=main,b=amd64
    origin archive.ubuntu.com
Pinned packages:
  gnucobol -> 2.2-5 with priority 500
```

_That is NOT what we want. _

For instance, checking libc-bin shows that it would have installed from the new non LTS distribution.

```
root@redshirt:~# apt-cache policy libc-bin
libc-bin:
  Installed: 2.27-3ubuntu1
  Candidate: 2.29-0ubuntu2
  Version table:
```

```

2.29-0ubuntu2 500
500 http://archive.ubuntu.com/ubuntu disco/main amd64 Packages
*** 2.27-3ubuntu1 500
500 http://archive.ubuntu.com/ubuntu bionic/main amd64 Packages
100 /var/lib/dpkg/status

```

Which would have likely wrecked havoc on system stability. *Good thing we are testing on a "redshirt"*

5.92.4 Stepwise Refinement

Guessing there was no default we modified the file to be more specific like this. (This was incorrect default is actually 500 for packages that aren't installed)

```

Package: *
Pin: release n=bionic*
Pin-Priority: 990

Package: *
Pin: release n=disco*
Pin-Priority: -10

Package: gncobol
Pin: release n=disco*
Pin-Priority: 500

```

Which at least fixes some of the issues.

```

root@redshirt:~# apt-cache policy
Package files:
100 /var/lib/dpkg/status
  release a=now
-10 http://security.ubuntu.com/ubuntu disco-security/universe amd64 Packages
  release v=19.04,o=Ubuntu,a=disco,n=disco,l=Ubuntu,c=universe,b=amd64
  origin security.ubuntu.com
...
-10 http://archive.ubuntu.com/ubuntu disco/restricted amd64 Packages
  release v=19.04,o=Ubuntu,a=disco,n=disco,l=Ubuntu,c=restricted,b=amd64
  origin archive.ubuntu.com
990 http://security.ubuntu.com/ubuntu bionic-security/multiverse amd64 Packages
  release v=18.04,o=Ubuntu,a=bionic-security,n=bionic,l=Ubuntu,c=multiverse,b=amd64
  origin security.ubuntu.com
...
990 http://archive.ubuntu.com/ubuntu bionic/main amd64 Packages
  release v=18.04,o=Ubuntu,a=bionic,n=bionic,l=Ubuntu,c=main,b=amd64
  origin archive.ubuntu.com
Pinned packages:
gncobol -> 2.2-5 with priority 500

```

Which gets us close. The stability is fixed.

```

root@redshirt:~# apt-cache policy libc-bin
libc-bin:
  Installed: 2.27-3ubuntu1
  Candidate: 2.27-3ubuntu1
  Version table:
   2.29-0ubuntu2 -10
   -10 http://archive.ubuntu.com/ubuntu disco/main amd64 Packages
*** 2.27-3ubuntu1 990
   990 http://archive.ubuntu.com/ubuntu bionic/main amd64 Packages
   100 /var/lib/dpkg/status

```

But the dependencies for the new package aren't.

```

root@redshirt:~# apt-get install --dry-run gncobol
Reading package lists... Done
Building dependency tree
Reading state information... Done
Some packages could not be installed. This may mean that you have
requested an impossible situation or if you are using the unstable
distribution that some required packages have not yet been created
or been moved out of Incoming.
The following information may help to resolve the situation:

The following packages have unmet dependencies:
gncobol : Depends: libcob4 but it is not installable
          Depends: libcob4-dev (= 2.2-5) but it is not installable
E: Unable to correct problems, you have held broken packages.
root@redshirt:~#

```

5.92.5 RTFM (*man apt_preferences*)

The apt preferences man pages explain a tiered priority system where ranges of numbers determine apt's behavior. Setting the priority for future packages to 100 allows missing packages to be installed.

```
root@redshirt:~# cat /etc/apt/preferences.d/gnucobol22.pref
Package: *
Pin: release n=bionic*
Pin-Priority: 990

Package: *
Pin: release n=disco*
Pin-Priority: 100

Package: gnucobol
Pin: release n=disco*
Pin-Priority: 600
```

Which works as we intended.

```
root@redshirt:~# apt-get update
Hit:1 http://archive.ubuntu.com/ubuntu bionic InRelease
Get:2 http://security.ubuntu.com/ubuntu bionic-security InRelease [88.7 kB]
Get:3 http://archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Get:4 http://security.ubuntu.com/ubuntu disco-security InRelease [97.5 kB]
Get:5 http://archive.ubuntu.com/ubuntu bionic-backports InRelease [74.6 kB]
Hit:6 http://archive.ubuntu.com/ubuntu disco InRelease
Fetched 350 kB in 2s (211 kB/s)
Reading package lists... Done
root@redshirt:~# apt-get install --dry-run gnucobol
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  binutils binutils-common binutils-x86-64-linux-gnu cpp cpp-7 gcc gcc-7 gcc-7-base libasan4 libatomic1 libbinutils libc-dev-bin libc6-dev libcc1-0 libcilkrts5 libcob4
  libcob4-dev libgcc-7-dev libgmp-dev libgmpxx4ldbl libgomp1 libisl19 libitm1 liblsan0 libmpc3 libmpx2 libncurses5-dev libncursesw6 libquadmath0 libtinfo-dev libtinfo6
  libtsan0
  libubsan0 linux-libc-dev manpages-dev
Suggested packages:
  binutils-doc cpp-doc gcc-7-locales gcc-multilib make autoconf automake libtool flex bison gdb gcc-doc gcc-7-multilib gcc-7-doc libgcc1-dbg libgomp1-dbg libitm1-dbg
  libatomic1-dbg libasan4-dbg liblsan0-dbg libubsan0-dbg libubsan0-dbg libmpx2-dbg libquadmath0-dbg glibc-doc gmp-doc libgmp10-doc libmpfr-dev ncurses-doc
The following NEW packages will be installed:
  binutils binutils-common binutils-x86-64-linux-gnu cpp cpp-7 gcc gcc-7 gcc-7-base gnucobol libasan4 libatomic1 libbinutils libc-dev-bin libc6-dev libcc1-0 libcilkrts5
  libcob4
  libcob4-dev libgcc-7-dev libgmp-dev libgmpxx4ldbl libgomp1 libisl19 libitm1 liblsan0 libmpc3 libmpx2 libncurses5-dev libncursesw6 libquadmath0 libtinfo-dev libtinfo6
  libtsan0
  libubsan0 linux-libc-dev manpages-dev
0 upgraded, 36 newly installed, 0 to remove and 3 not upgraded.
Inst libtinfo6 (6.1+20181013-2ubuntu2 Ubuntu:19.04/disco [amd64])
Inst libncursesw6 (6.1+20181013-2ubuntu2 Ubuntu:19.04/disco [amd64])
Inst binutils-common (2.30-21ubuntu1~18.04.2 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libbinutils (2.30-21ubuntu1~18.04.2 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst binutils-x86-64-linux-gnu (2.30-21ubuntu1~18.04.2 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst binutils (2.30-21ubuntu1~18.04.2 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst gcc-7-base (7.4.0-1ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libisl19 (0.19-1 Ubuntu:18.04/bionic [amd64])
Inst libmpc3 (1.1.0-1 Ubuntu:18.04/bionic, Ubuntu:19.04/disco [amd64])
Inst cpp-7 (7.4.0-1ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst cpp (4:7.4.0-1ubuntu2.3 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libcc1-0 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libgomp1 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libitm1 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libatomic1 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libasan4 (7.4.0-1ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst liblsan0 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libtsan0 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libubsan0 (7.4.0-1ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libcilkrts5 (7.4.0-1ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libmpx2 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libquadmath0 (8.3.0-6ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libgcc-7-dev (7.4.0-1ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst gcc-7 (7.4.0-1ubuntu1~18.04.1 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst gcc (4:7.4.0-1ubuntu2.3 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libc-dev-bin (2.27-3ubuntu1 Ubuntu:18.04/bionic [amd64])
Inst linux-libc-dev (4.15.0-54.58 Ubuntu:18.04/bionic-updates, Ubuntu:18.04/bionic-security [amd64])
Inst libc6-dev (2.27-3ubuntu1 Ubuntu:18.04/bionic [amd64])
Inst libgmpxx4ldbl (2:6.1.2+dfsg-2 Ubuntu:18.04/bionic [amd64])
Inst libgmp-dev (2:6.1.2+dfsg-2 Ubuntu:18.04/bionic [amd64])
Inst libtinfo-dev (6.1-1ubuntu1.18.04 Ubuntu:18.04/bionic-updates [amd64])
Inst libncurses5-dev (6.1-1ubuntu1.18.04 Ubuntu:18.04/bionic-updates [amd64])
Inst manpages-dev (4.15-1 Ubuntu:18.04/bionic [all])
Inst libcob4 (2.2-5 Ubuntu:19.04/disco [amd64])
Inst libcob4-dev (2.2-5 Ubuntu:19.04/disco [amd64])
Inst gnucobol (2.2-5 Ubuntu:19.04/disco [amd64])
Conf libtinfo6 (6.1+20181013-2ubuntu2 Ubuntu:19.04/disco [amd64])
...
Conf gnucobol (2.2-5 Ubuntu:19.04/disco [amd64])
root@redshirt:~#
```

5.92.6 References

- <https://medium.com/@george.shuklin/how-to-install-packages-from-a-newer-distribution-without-installing-unwanted-6584fa93208f>
- <https://askubuntu.com/questions/49609/how-do-i-add-the-proposed-repository>

5.93 Install Ansible

All centralized maintainance should be initiated from kb2018

```
root@kb2018:~# apt-get install ansible
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  ieee-data python-asn1crypto python-certifi python-cffi-backend python-chardet python-cryptography python-enum34 python-httplib2 python-idna python-ipaddress
  python-jinja2 python-jmespath python-kerberos python-libcloud python-lockfile python-markupsafe python-netaddr python-openssl python-paramiko python-pkg-resources
  python-pyasn1 python-requests python-selinux python-simplejson python-six python-urllib3 python-xlrd python-yaml
Suggested packages:
  cowsay sshpass python-cryptography-doc python-cryptography-vectors python-enum34-doc python-jinja2-doc python-lockfile-doc ipython python-netaddr-docs
  python-openssl-doc python-openssl-dbg python-gssapi python-setuptools python-socks python-ntlm
...
root@kb2018:~#
```

5.93.1 Install python to all containers

```
root@kb2018:~# for h in `lxc list local: -c n --format csv`;do echo $h;lxc exec local:$h -- apt-get install -y python; done
...
root@kb2018:~# for h in `lxc list bs2020: -c n --format csv`;do echo $h;lxc exec bs2020:$h -- apt-get install -y python; done
...
```

5.93.2 Seed /etc/ansible/hosts

localhost (kb2018)

Adding the entry for the localhost is simple

```
root@kb2018:~# nano /etc/ansible/hosts

[pets:children]
servers
containers

[servers]
kb2018  ansible_connection=local
..
root@kb2018:~#
```

local containers

entries for local containers is equally straightforward.

```
hostname ansible_connection=lxd
```

Which we can generate using lxc list and awk

```
root@kb2018:~# lxc list -c n --format=csv local:|awk '{print $1,"ansible_connection=lxd";}'>>/etc/ansible/hosts
```

containers on remote host

Containers on the remote host (bs2020) require an additional parameter

```
remotecontainer  ansible_connection=lxd ansible_host=remotehost:remotecontainer
```

Which we again generate using lxc list and awk

```
root@kb2018:~# lxc list -c n --format=csv bs2020:|awk '{print $1," ansible_connection=lxd ansible_host=bs2020:"$1;}'>>/etc/ansible/hosts
```


5.93.3 adding access to bs2020 (via ssh to unprivileged account)

Our current security model expressly forbids direct access to all root accounts, users must connect using an ssh key and escalate using their password.

To control a remote server from ansible user (root@kb2018) we:

Create a sudo user for our ansible host

```
root@bs2020:~# useradd kb2018 -c"Governor Kate Brown" -m -g sudo
root@bs2020:~# passwd kb2018
... remember this one for later ...
```

Restrict ssh access to that account to the ip of that particular host.

```
root@bs2020:~# nano /etc/ssh/sshd_config
...
PermitRootLogin no
....
DenyUsers kb2018@"!192.168.31.159,*"
...
root@bs2020:~# service ssh restart
```

Generate key for our ansible user (root@kb2018)

```
haifisch:~ don$ ssh -p22222 feurig@bs2020.suspectdevices.com
...
Last login: Mon Feb 25 18:56:59 2019 from 97.115.103.251
feurig@kb2018:~$ sudo bash
[sudo] password for feurig:
root@kb2018:~# ssh-key
ssh-keygen ssh-keyscan
root@kb2018:~# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
...
... add ssh key to kb2018@bs2020:~/.ssh/authorized_keys ...
...
root@kb2018:~# ssh kb2018@bs2020.suspectdevices.com
```

5.93.4 Testing connectivity.

At this point we can add the remote server to ansible's inventory and check the connectivity.

```
bs2020  ansible_connection=ssh ansible_ssh_user=kb2018
```

note kb2018 is the localhost, ernest24jan19 (stopped) and douglas are local containers, bs2020 is a remote host and teddy is a container that it hosts

```
root@kb2018:~# ansible pets -m ping
kb2018 | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
ernest24jan19 | UNREACHABLE! => {
  "changed": false,
  "msg": "... , exited with result 1",
  "unreachable": true
}
...
douglas | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
...
bs2020 | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
...
teddy | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
```

```
}
root@kb2018:~#
```

However we cannot run privileged commands on our remote host.

```
root@kb2018:~# ansible servers -m apt -a "force_apt_get=yes upgrade=yes update_cache=yes autoremove=yes"
bs2020 | FAILED! => {
  "changed": false,
  "msg": "Failed to lock apt for exclusive operation"
}
kb2018 | SUCCESS => {
```

We can fix this by telling ansible to escalate using our user and password

```
bs2020  ansible_host=bs2020.suspectdevices.com ansible_user=kb2018ansible_become=yes ansible_become_user=root ansible_become_pass=my_super_secret_password
```

And we can see that this works. Next we encrypt the password using ansible's vault feature and moving the username and password to the host_vars file.

```
feurig@kb2018:~$ grep 'bs2020 ' /etc/ansible/hosts
bs2020  ansible_host=bs2020.suspectdevices.com ansible_user='{{ bs2020_unprivileged_user }}' ansible_become=yes ansible_become_user=root
ansible_become_pass='{{ bs2020_become_pass }}'
root@kb2018:~# ansible servers -m apt -a "force_apt_get=yes upgrade=yes update_cache=yes autoremove=yes"
kb2018 | SUCCESS => {
```

... *WIP: you are here* ... Create and protect vault password file

```
root@kb2018:~# openssl rand -base64 2048 > /root/.vault_passwd
root@kb2018:~# chmod 600 /root/.vault_passwd
```

Add password file to ansible.cfg

```
root@kb2018:~# nano /etc/ansible/ansible.cfg
...
# If set, configures the path to the Vault password file as an alternative to
# specifying --vault-password-file on the command line.
#vault_password_file = /path/to/vault_password_file
vault_password_file=/root/.vault_passwd
...
```

Encrypt sudo password

```
root@kb2018:~# ansible-vault encrypt_string 'mybigsecret' --name 'kb2018_become_pass'
kb2018_become_pass: !vault |
    $ANSIBLE_VAULT;1.1;AES256
    663 .... 462
Encryption successful
```

Add user and encrypted password to /etc/ansible/host_vars/bs2020.yml

```
root@kb2018:~# mkdir /etc/ansible/host_vars
root@kb2018:~# nano /etc/ansible/host_vars/bs2020.yml
bs2020_unprivileged_user: kb2018
bs2020_become_pass: !vault |
    $ANSIBLE_VAULT;1.1;AES256
    663 .... 462
```

Add variables to inventory

```
[pets:children]
servers
containers

[servers]
kb2018  ansible_connection=local
bs2020  ansible_host=bs2020.suspectdevices.com ansible_user '{{ bs2020_unprivileged_user }}' ansible_become=yes ansible_become_user=root
ansible_become_pass '{{ bs2020_become_pass }}'
#bs2020  ansible_connection=ssh ansible_ssh_user=kb2018

[containers:children]
local-containers
remote-containers

[local-containers]
douglas ansible_connection=lxd
...
[remote-containers]
...
goethe  ansible_connection=lxd ansible_host=bs2020:goethe
```

And now we can treat all of our pets with the same love and affection.

```
root@kb2018:~# ansible pets -m apt -a "force_apt_get=yes upgrade=yes update_cache=yes autoremove=yes"
```

5.93.5 References/Linkdump

- <https://stackoverflow.com/questions/37297249/how-to-store-ansible-become-pass-in-a-vault-and-how-to-use-it>
- https://docs.ansible.com/ansible/latest/user_guide/vault.html#id6

5.94 Squid Caching Server

(... explain what we want to get our of squid ... Basic proxy ... proxy forwarded to remote proxy reverse proxy ... more words here....)

5.94.1 Why is this taking so long ???

"I installed squid3 (on Ubuntu), but looking at the configuration file, I am lost. I tried googling but it looks all too complicated." -- stack overflow user

Squid is a monster. Years of development and added features have created a configuration file that has 8000 lines of comments and 20 actual lines of configuration which need to be modified for it to work at all.

5.94.2 basic proxy configuration

The configuration below can be found on the squid containers on both basement servers (Joey and DeeDee). To use this server set your web browser proxy to the http_port in the configuration file (3128).

```
acl SSL_ports port 443
acl Safe_ports port 80      # http
acl Safe_ports port 21      # ftp
acl Safe_ports port 443     # https
acl Safe_ports port 70      # gopher
acl Safe_ports port 210     # wais
acl Safe_ports port 1025-65535 # unregistered ports
acl Safe_ports port 280     # http-mgmt
acl Safe_ports port 488     # gss-http
acl Safe_ports port 591     # filemaker
acl Safe_ports port 777     # multiling http
acl CONNECT method CONNECT
http_access deny !Safe_ports
http_access deny CONNECT !SSL_ports
http_access allow localhost manager
http_access deny manager
http_access allow localhost
acl my_internal_net src 192.168.0.0/24
http_access allow my_internal_net
http_port 3128
coredump_dir /var/spool/squid
refresh_pattern ^ftp:      1440  20%  10080
refresh_pattern ^gopher:  1440  0%   1440
refresh_pattern -i (/cgi-bin/|\?) 0    0%   0
refresh_pattern (Release|Packages|.gz)*$ 0    20%  2880
refresh_pattern .          0    20%  4320
```

5.94.3 reverse proxy configuration

Work in progress.

minimal (no ssl) configuration

With an ssh tunnel set from our local static web server coming onto the server (jules.suspectdevices.com) at port 8085 we tell squid to route all traffic on port 80 to the server on the other end of the tunnel.

```
debug_options ALL,2 28,9
http_port 80 accel no-vhost defaultsite=jules.suspectdevices.com
cache_peer 127.0.0.1 parent 8085 0 no-query originserver name=corbin
acl theworld src all
acl our_sites dstdomain all
http_access allow theworld
http_access allow our_sites
cache_peer_access corbin allow our_sites
cache_peer_access corbin allow theworld
http_access deny all
```

Secure reverse proxy

... working on it ...

From: squid example configurations

```
https_port 443 accel defaultsite=jules.suspectdevices.com \
cert=/etc/ssl/certs/ssl-cert-snakeoil.pem \
key=/etc/ssl/private/ssl-cert-snakeoil.key

# First (HTTP) peer
cache_peer 127.0.0.1 parent 8086 0 no-query originserver login=PASS name=lilly

acl pdx dstdomain jules.suspectdevices.com
cache_peer_access lilly allow pdx
http_access allow pdx

# Security block for non-hosted sites
http_access deny all
```

From: <https://serverfault.com/questions/735535/squid-reverse-proxy-redirect-rewrite-http-to-https>

```
acl PORT80 myport 80
http_access deny PORT80 pdx
deny_info 301:https://foo.server.com%R pdx
```

5.94.4 Dual Proxy Configuration

TODO: Having a local proxy combined with a "pihole" dns server seems to improve browsing performance considerably. Adding a second proxy upstream would allow less location based garbage as well. [TaskDualProxyConfiguration Dual Proxy Configuration Notes]

5.94.5 Preliminary Linkdump

- <http://cosmolinux.no-ip.org/raconetlinux/html/17-squid.html>
- https://wiki.squid-cache.org/SquidFaq/ConfiguringSquid#Before_you_start_configuring
- <https://www.tekyhost.com/squid-proxy-squid-caching-and-filtering-proxy/>
- <https://www.rootusers.com/configure-squid-proxy-to-forward-to-a-parent-proxy/>
- <https://wiki.squid-cache.org/Features/CacheHierarchy>
- <https://www.tecmint.com/install-squid-in-ubuntu/>
- <http://www.squidguard.org/about.html>
- https://wiki.alpinelinux.org/wiki/Setting_up_Transparent_Squid_Proxy

5.95 Task: Split ZF Mirror

We need to reduce the size of the zfs pool on the two 600G disks on bs2020 and use the space for backups. To do this we need to split the mirror and use the freed disk to create a new partition. Move that into place and then move the old data onto the smaller partition before finally repartitioning the remaining disk and mirror both new partitions.

Move running containers to kb2018

Get existing disk info.

```
root@bs2020:~# zpool status -L devel
pool: devel
state: ONLINE
scan: resilvered 132G in 2h2m with 0 errors on Thu Apr  4 00:31:41 2019
config:

    NAME        STATE      READ WRITE CKSUM
    devel        ONLINE     0   0   0
      mirror-0   ONLINE     0   0   0
        sdc      ONLINE     0   0   0
        sdd      ONLINE     0   0   0

errors: No known data errors
root@bs2020:~# zpool status  devel
pool: devel
state: ONLINE
scan: resilvered 132G in 2h2m with 0 errors on Thu Apr  4 00:41:41 2019
config:

    NAME                                STATE      READ WRITE CKSUM
    devel                                ONLINE     0   0   0
      mirror-0                           ONLINE     0   0   0
        scsi-350000c0f022fd4c8           ONLINE     0   0   0
        scsi-35000c50047d0926f           ONLINE     0   0   0

errors: No known data errors
```

Split devel mirror.

```
root@bs2020:~# zpool split -R /newdevel devel newdevel
root@bs2020:~# zpool status
pool: devel
state: ONLINE
scan: scrub repaired 0B in 0h38m with 0 errors on Sun Mar 10 01:02:05 2019
config:

    NAME                                STATE      READ WRITE CKSUM
    devel                                ONLINE     0   0   0
      scsi-35000c50047d0926f             ONLINE     0   0   0

errors: No known data errors

pool: infra
state: ONLINE
scan: scrub repaired 0B in 0h1m with 0 errors on Sun Mar 10 00:25:17 2019
config:

    NAME                                STATE      READ WRITE CKSUM
    infra                                ONLINE     0   0   0
      mirror-0                           ONLINE     0   0   0
        scsi-35000cca00b33a264           ONLINE     0   0   0
        scsi-350000395a8336d34           ONLINE     0   0   0

errors: No known data errors

pool: newdevel
state: ONLINE
scan: scrub repaired 0B in 0h38m with 0 errors on Sun Mar 10 01:02:05 2019
config:

    NAME                                STATE      READ WRITE CKSUM
    newdevel                             ONLINE     0   0   0
      scsi-350000c0f022fd4c8             ONLINE     0   0   0
```

Wipe and partition newly freed disk

```
root@bs2020:~# zpool destroy newdevel
root@bs2020:~# parted /dev/sdc
GNU Parted 3.2
Using /dev/sdc
(parted) mklabel gpt
Warning: The existing disk label on /dev/sdc will be destroyed and all data on this disk will be lost. Do you want to continue?
```

```
Yes/No? yes
(parted) mkpart
Partition name? []? images
File system type? [ext2]? zfs
Start? 0%
End? 50%
(parted) mkpart
Partition name? []? devel
File system type? [ext2]? zfs
Start? 50%
End? 100%
(parted) print
Model: WD WD6001BKHG (scsi)
Disk /dev/sdc: 600GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB   300GB   300GB    zfs          images
  2      300GB   600GB   300GB    zfs          devel

(parted) quit
```

Build new zfs partition for /var/lib/lxd/images and move the old data

```
root@bs2020:~# systemctl stop lxd
root@bs2020:~# zpool create lxd-images scsi-350000c0f022fd4c8-part1 -m/var/lib/lxd/images
root@bs2020:~# df -k
Filesystem      1K-blocks      Used Available Use% Mounted on
udev             49458232         0  49458232   0% /dev
tmpfs            9897784       1488   9896296   1% /run
/dev/sdg2       138930656  63784172  68019548  49% /
tmpfs            49488916         0  49488916   0% /dev/shm
tmpfs            5120          0       5120   0% /run/lock
tmpfs            49488916         0  49488916   0% /sys/fs/cgroup
/dev/loop0       53376        53376         0 100% /snap/lxd/10234
/dev/loop1       91392        91392         0 100% /snap/core/6673
/dev/loop2       55168        55168         0 100% /snap/lxd/10343
/dev/loop3       93312        93312         0 100% /snap/core/6531
/dev/loop4       53376        53376         0 100% /snap/lxd/10289
/dev/loop5       93184        93184         0 100% /snap/core/6405
/dev/sdg1        523248        6164   517084    2% /boot/efi
/dev/sda1       480589544  47764176  408389708 11% /archive
tmpfs            9897780         0   9897780   0% /run/user/1000
lxd-images      282394496         0  282394496   0% /var/lib/lxd/images
root@bs2020:~# mv /var/lib/lxd/images.tmp/* /var/lib/lxd/images/
df -k
root@bs2020:~# df -k
Filesystem      1K-blocks      Used Available Use% Mounted on
udev             49458232         0  49458232   0% /dev
tmpfs            9897784       1496   9896288   1% /run
/dev/sdg2       138930656  15992764 115810956 13% /
tmpfs            49488916         0  49488916   0% /dev/shm
tmpfs            5120          0       5120   0% /run/lock
tmpfs            49488916         0  49488916   0% /sys/fs/cgroup
/dev/loop0       53376        53376         0 100% /snap/lxd/10234
/dev/loop1       91392        91392         0 100% /snap/core/6673
/dev/loop2       55168        55168         0 100% /snap/lxd/10343
/dev/loop3       93312        93312         0 100% /snap/core/6531
/dev/loop4       53376        53376         0 100% /snap/lxd/10289
/dev/loop5       93184        93184         0 100% /snap/core/6405
/dev/sdg1        523248        6164   517084    2% /boot/efi
/dev/sda1       480589544  47764176  408389708 11% /archive
tmpfs            9897780         0   9897780   0% /run/user/1000
lxd-images      282392832  47751552 234641280 17% /var/lib/lxd/images
root@bs2020:~# systemctl start lxd
```

Build new devel pool on second partition and move data to smaller partition..

```
root@bs2020:~# zpool create devels scsi-350000c0f022fd4c8-part2
root@bs2020:~# zfs snapshot -r devel@fullbackup
root@bs2020:~# zfs send -R devel@fullbackup | pv | zfs receive -vFdu devels
....
```

Offline old devel pool replace with new

```
root@bs2020:~# systemctl stop lxd
root@bs2020:~# zpool export devels
root@bs2020:~# zpool destroy devel
root@bs2020:~# zpool import devels devel
root@bs2020:~# zpool status
pool: devel
state: ONLINE
scan: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    devel               ONLINE                0     0     0
```

```
scsi-350000c0f022fd4c8-part2 ONLINE 0 0 0

errors: No known data errors
...
root@bs2020:~# systemctl start lxd
```

Repartition remaining disk.

```
root@bs2020:~# parted /dev/sdd
GNU Parted 3.2
Using /dev/sdd
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) print
Model: SEAGATE ST9600205SS (scsi)
Disk /dev/sdd: 600GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End    Size  File system  Name              Flags
 1      1049kB 600GB  600GB  zfs          zfs-f77cf42a401f4fa0
 9      600GB  600GB  8389kB

(parted) mklabel
New disk label type? gpt
Warning: The existing disk label on /dev/sdd will be destroyed and all data on this disk will be lost. Do you want to continue?
Yes/No? yes
(parted) mkpart
Partition name? []? images
File system type? [ext2]? zfs
Start? 0%
End? 50%
(parted) mkpart
Partition name? []? devel
File system type? [ext2]? zfs
Start? 50%
End? 100%
(parted) print
Model: SEAGATE ST9600205SS (scsi)
Disk /dev/sdd: 600GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End    Size  File system  Name      Flags
 1      1049kB 300GB  300GB  zfs          images
 2      300GB  600GB  300GB  zfs          devel

(parted) quit
Information: You may need to update /etc/fstab.
```

Add new partitions as mirrors

```
root@bs2020:~# zpool attach lxd-images scsi-350000c0f022fd4c8-part1 scsi-35000c50047d0926f-part1
root@bs2020:~# zpool attach devel scsi-350000c0f022fd4c8-part2 scsi-35000c50047d0926f-part2
root@bs2020:~# zpool status
pool: devel
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
        continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Fri Apr  5 22:39:37 2019
      50.0M scanned out of 132G at 2.94M/s, 12h42m to go
      48.6M resilvered, 0.04% done
config:

        NAME                                STATE    READ WRITE CKSUM
        devel                                ONLINE      0   0   0
          mirror-0                            ONLINE      0   0   0
            scsi-350000c0f022fd4c8-part2      ONLINE      0   0   0
            scsi-35000c50047d0926f-part2      ONLINE      0   0   0 (resilvering)

errors: No known data errors

pool: infra
state: ONLINE
scan: scrub repaired 0B in 0h1m with 0 errors on Sun Mar 10 00:25:17 2019
config:

        NAME                                STATE    READ WRITE CKSUM
        infra                                ONLINE      0   0   0
          mirror-0                            ONLINE      0   0   0
            scsi-35000cca00b33a264            ONLINE      0   0   0
            scsi-350000395a8336d34            ONLINE      0   0   0

errors: No known data errors

pool: lxd-images
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
        continue to function, possibly in a degraded state.
```



```
action: Wait for the resilver to complete.
scan: resilver in progress since Fri Apr  5 22:39:09 2019
      3.70G scanned out of 45.6G at 84.2M/s, 0h8m to go
      3.70G resilvered, 8.11% done
config:

NAME                STATE  READ WRITE CKSUM
lxd-images           ONLINE    0     0     0
mirror-0             ONLINE    0     0     0
scsi-350000c0f022fd4c8-part1 ONLINE    0     0     0
scsi-35000c50047d0926f-part1 ONLINE    0     0     0 (resilvering)

errors: No known data errors
root@bs2020:~#
```

Move containers back to bs2020

Continue work on backup scripts.

5.95.1 References

<https://github.com/lxc/lxd/issues/4984>

5.96 Task: ZFS Disk Replacement

The process of replacing mirrored zfs disks is fairly simple. The changes are done by zpool attach and detach.

```
zpool detach <pool> <disk-id>
zpool attach <pool> <disk-id-to-mirror> <disk-id-mirrored-to>
```

The heavy lifting is done by zfs itself.

5.96.1 process

PREP

- use the [pdf article link](#) to print this before going down
- If possible pre wipe and check the disks on a separate linux machine (*note: /dev/sdf is an placeholder for the disk mounted on that system*)

```
root@homebox:~# wipefs -af --backup /dev/sdf /dev/sdf: 8 bytes were erased at offset 0x00000200 (gpt): 45 46 49 20 50 41
52 54 /dev/sdf: 8 bytes were erased at offset 0x222ee64e00 (gpt): 45 46 49 20 50 41 52 54 /dev/sdf: 2 bytes were erased at
offset 0x000001fe (PMBR): 55 aa /dev/sdc: calling ioctl to re-read partition table: Success root@homebox:~# fdisk /dev/sdf
.... Command (m for help): g
```

Created a new GPT disklabel (GUID: EBC5A0C9-E871-544F-A8EA-E31FCA655F9C).

Command (m for help): w The partition table has been altered. Calling ioctl() to re-read partition table. Syncing disks.

```
root@homebox:~# badblocks /dev/sdf ....
```

- insure that you can ssh into the box

On site

The following assumes you have escalated to root privileges (sudo bash), in this case we are replacing /dev/sdc and /dev/sdd in the pool named 'level'

- check for the correct disk. The following should cause the disk to light up\ (*C when you have identified the disk. Careful with the if/of here*).

```
root@bs2020:~# dd if=/dev/sdc of=/dev/null
```

- find the disk in the pool.

```
root@bs2020:~# zpool status pool: devel state: ONLINE scan: resilvered 9.95G in 0h4m with 0 errors on Sat Nov 10
22:00:41 2018 config:
```

NAME	STATE	READ	WRITE	CKSUM
devel	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
scsi-35000c50054fee503	ONLINE	0	0	0
scsi-35000c5005501b45b	ONLINE	0	0	0

errors: No known data errors

```
... root@bs2020:~# ls -ls /dev/disk/by-id/|grep scsi|grep -v "-part" 0 lrwxrwxrwx 1 root root 9 Nov 10 21:22
scsi-350000395a8336d34 -> ../../sde 0 lrwxrwxrwx 1 root root 9 Nov 10 21:22 scsi-35000c50054fee503 -> ../../sdd 0
lrwxrwxrwx 1 root root 9 Nov 10 21:56 scsi-35000c5005501b45b -> ../../sdc 0 lrwxrwxrwx 1 root root 9 Nov 10 21:22
scsi-35000cca00b33a264 -> ../../sdf 0 lrwxrwxrwx 1 root root 9 Nov 10 21:22 scsi-3600508e00000000069cf3977618f1408
-> ../../sdg root@bs2020:~#
```

We notice above that the disk we are looking for is scsi-35000c5005501b45b

- detach the disk from the pool.

```
root@bs2020:~# zpool detach devel scsi-35000c5005501b45b root@bs2020:~# zpool status pool: devel state: ONLINE
scan: resilvered 9.95G in 0h4m with 0 errors on Sat Nov 10 22:00:41 2018 config:
```

NAME	STATE	READ	WRITE	CKSUM
devel	ONLINE	0	0	0
scsi-35000c50054fee503	ONLINE	0	0	0

errors: No known data errors

pool: infra ... root@bs2020:~#

- even if expanding the disk size insure that auto expand is off.

```
root@bs2020:~# zpool set autoexpand=off devel
```

- Swap out the old disk with the new one.
- find the new disk's id.

```
root@bs2020:~# partprobe root@bs2020:~# ls -ls /dev/disk/by-id/|grep sdc 0 lrwxrwxrwx 1 root root 9 Nov 10 21:56 scsi-
xxxxxxxxxxxxxxxxxx -> ../sdc ...
0 lrwxrwxrwx 1 root root 9 Nov 10 21:56 xxx-xxxxxxxxxxxxxxxxxx -> ../sdc
```

- *If the drive id does not change reboot the server*
- attach the new disk to the zfs pool (*scsi-xxxxxxxxxxxxxxxxxx is the new id from the above step*)

```
root@bs2020:~# zpool attach devel scsi-35000c50054fee503 scsi-xxxxxxxxxxxxxxxxxx
```

- wait for pool to resilver

```
root@bs2020:~# zpool status pool: devel state: ONLINE status: One or more devices is currently being resilvered. The pool
will continue to function, possibly in a degraded state. action: Wait for the resilver to complete. scan: resilver in progress since
Sat Nov 10 21:56:04 2018 8.54G scanned out of 9.95G at 35.5M/s, 0h0m to go 8.54G resilvered, 85.85% done config:
```

NAME	STATE	READ	WRITE	CKSUM
devel	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
scsi-35000c50054fee503	ONLINE	0	0	0
scsi-35000c5005501b45b	ONLINE	0	0	0 (resilvering)

errors: No known data errors

pool: infra ...

```
root@bs2020:~# zpool status .... repeat until finished resilvering .... root@bs2020:~# zpool status pool: devel state: ONLINE
scan: scrub repaired 0B in 0h4m with 0 errors on Sat Nov 10 21:58:04 2018 config:
```

NAME	STATE	READ	WRITE	CKSUM
devel	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
scsi-35000cca00b33a264	ONLINE	0	0	0
scsi-350000395a8336d34	ONLINE	0	0	0

errors: No known data errors ...

- if expanding disk check for new size and if not expand it

5.97 zfs list and check for larger disk pool

- repeat process for disk in bay below (we already know its old id from above).

```
root@bs2020:~# dd if=/dev/sdd of=/dev/null root@bs2020:~# zpool detach devel scsi-35000c50054fee503
... swap disks ... root@bs2020:~# partprobe root@bs2020:~# ls -ls /dev/disk/by-id/|grep sdd 0 lrwxrwxrwx 1 root root 9 Nov
10 21:56 scsi-yyy-yyyyyyyyyyyyyy-> ../sdd ... reboot if necessary ... root@bs2020:~# wipefs -a /dev/sdd ... root@bs2020:~#
fdisk /dev/sdd ... root@bs2020:~# zpool attach devel scsi-xxxxxxxxxxxxxxxxxx scsi-yyy-yyyyyyyyyyyyyy ... wait for resilver...
```

- use the process below to grow disks to new size

5.98 zpool set autoexpand=on devel

5.99 zpool online -e devel scsi-xxxxxxxxxxxxxxxxxxxxxx

5.100 zpool online -e devel scsi-yyyyyyyyyyyyyyyyyyyy

5.101 zpool set autoexpand=off devel

references

- <https://tomasz.korwel.net/2014/01/03/growing-zfs-pool/>
- <https://jsosic.wordpress.com/2013/01/01/expanding-zfs-zpool-raid/>
- <https://serverfault.com/questions/5336/how-do-i-make-linux-recognize-a-new-sata-dev-sda-drive-i-hot-swapped-in-without>

5.102 Ubuntu18.04Notes

5.102.1 Netplan / Networkd

Given the success of systemd the kids decided that they needed to rewrite the networking core using a yaml file under /etc/netplan/ and various "renderers". If it all gets too much you can replace it with the legacy system ifupdown and continue to edit /etc/network/interfaces, etc.

```
apt-get install ifupdown
```

Otherwise read the notes to follow.

See: [Netplan Documentation \(https://netplan.io/\)](https://netplan.io/)

Static Networking with Netplan

Assuming that your cloud configuration does not overwrite it the following file produces a static ip.

```
oot@phillip:~# cat /etc/netplan/50-cloud-init.yaml
network:
  version: 2
  ethernets:
    eth0:
      dhcp4: no
      addresses: [198.202.31.223/25]
      gateway4: 198.202.31.129
      nameservers:
        search: [suspectdevices.com fromhell.com vpn]
        addresses: [198.202.31.141]
```

Bridge Networking with Netplan

```
root@annie:~# nano /etc/netplan/01-netcfg.yaml
network:
  version: 2
  renderer: networkd
  ethernets:
    ens6:
      dhcp4: true
      dhcp6: no
    enpls0:
      dhcp4: no
      dhcp6: no
  bridges:
    br0:
      dhcp4: no
      dhcp6: no
      addresses:
        - 192.168.0.66/24
      gateway4: 192.168.0.1
      nameservers:
        addresses:
          - 192.168.0.1
          - 198.202.31.141
      interfaces:
        - enpls0
root@annie:~# netplan apply
root@annie:~# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
...
2: enpls0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master br0 state UP group default qlen 1000
    link/ether 78:e7:d1:c3:ef:9e brd ff:ff:ff:ff:ff:ff
3: ens6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether 00:14:d1:25:2b:bc brd ff:ff:ff:ff:ff:ff
    inet 192.168.2.66/24 brd 192.168.2.255 scope global dynamic ens6
        valid_lft 43163sec preferred_lft 43163sec
    inet6 fd5b:alad:aeeb::fd0/128 scope global noprefixroute
...
6: br0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether aa:18:c9:5a:76:d6 brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.66/24 brd 192.168.0.255 scope global br0
        valid_lft forever preferred_lft forever
    inet6 fe80::a818:c9ff:fe5a:76d6/64 scope link
        valid_lft forever preferred_lft forever
root@annie:~# brctl show
```

```

bridge name bridge id      STP enabled interfaces
br0      8000.aa18c95a76d6  no      enpls0
root@annie:~#

```

5.102.2 And it works for anonymous bridges EXCEPT FOR THE BUG

Basically if no address is given for a bridge netplan fails to tell systemd to up the interface anyway and the bridges do not come up.

```

root@bs2020:~# nano /etc/netplan/50-cloud-init.yaml
# This file is generated from information provided by
# the datasource.  Changes to it will not persist across an instance.
# To disable cloud-init's network configuration capabilities, write a file
# /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg with the following:
# network: {config: disabled}
network:
  version: 2
  renderer: networkd
  ethernets:
    eno1:
      dhcp4: no
      addresses: [192.168.31.158/24]
      gateway4: 192.168.31.1
      nameservers:
        search: [suspectdevices.com fromhell.com vpn]
        addresses: [198.202.31.141]
    eno2:
      dhcp4: no
      optional: true
    eno3:
      dhcp4: no
    eno4:
      dhcp4: no
  bridges:
    br0:
      dhcp4: no
      dhcp6: no
      interfaces:
        - eno4
    br1:
      dhcp4: no
      dhcp6: no
      interfaces:
        - eno3

root@bs2020:~# nano /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg
network: {config: disabled}
root@bs2020:~# netplan apply

```

- so you have to create the scripts until they fix this.

```

root@bs2020:~# nano /etc/systemd/network/br0.network [Match] Name=br0
[Network] LinkLocalAddressing=no IPv6AcceptRA=no

root@bs2020:~# nano /etc/systemd/network/br1.network [Match] Name=br1
[Network] LinkLocalAddressing=no IPv6AcceptRA=no

```

<https://bugs.launchpad.net/ubuntu/+source/nplan/+bug/1736975> <http://djanotes.blogspot.com/2018/04/anonymous-bridges-in-netplan.html>

Freaking Cloud init

Need to figure out how much damage is done here...

Starting with the hostname. The hostname is now handled by a new command and /etc/cloud/cloud.config needs to be modified to preserve the hostname across boots.

```

feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# hostnamectl set-hostname bs2020
root@bs2020:~# nano /etc/cloud/cloud.cfg
....
# This will cause the set+update hostname module to not operate (if true)
preserve_hostname: true
...
root@bs2020:~# reboot

```

Install the root users .

One would like for the installer to give you some options for installing the admin team but we just paste the hash from one of the other machines into the shadow password file and copy the home directories for their ssh keys. see [wiki:kb2018InstallBashHistory](#)

```
( .. tired of winning .... write up later... )
```

Install zfs

```
root@bs2020:~# apt-get install nfs-kernel-server samba-common-bin zfsutils-linux
```

- create zfs pools using lxd init.
- make servers available to each other.
- configure outgoing mail.
- install apticron

Link Dump

- <https://netplan.io/examples>
- <https://websiteforstudents.com/configure-static-ip-addresses-on-ubuntu-18-04-beta/>
- <https://askubuntu.com/questions/1054350/netplan-bridge-for-kvm-on-ubuntu-server-18-04-with-static-ips> <https://stackoverflow.com/questions/33377916/migrating-lxc-to-lxd>

5.103 Ubuntu LTS Email Server Setup

This document assumes that you have set up a debian 9 or ubuntu LTS server(/container) set up and that postfix/email has been set up using tasksel.

5.103.1 Dovecot (imap server) and Postfix (mail server)

configure dovecot to use self signed ssl cert created by postfix.

```
root@naomi:/etc/postfix# cd ../dovecot/conf.d/
root@naomi:/etc/dovecot/conf.d# nano 10-ssl.conf
##
## SSL settings
##

# SSL/TLS support: yes, no, required. <doc/wiki/SSL.txt>
ssl = yes

# PEM encoded X.509 SSL/TLS certificate and private key. They're opened before
# dropping root privileges, so keep the key file unreadable by anyone but
# root. Included doc/mkcert.sh can be used to easily generate self-signed
# certificate, just make sure to update the domains in dovecot-openssl.cnf
ssl_cert = </etc/ssl/certs/ssl-cert-snakeoil.pem
ssl_key = </etc/ssl/private/ssl-cert-snakeoil.key
#ssl_cert = </etc/dovecot/dovecot.pem
#ssl_key = </etc/dovecot/private/dovecot.pem
```

Also set mailbox format to Maildir or all of your legacy data will be hosed.

```
root@naomi:/etc/dovecot/conf.d# nano 10-mail.conf
mail_location = maildir:~/Maildir
...
```

Notice issues with sending mail using ssl/tls

```
don@bob2:~$ openssl s_client -connect mail.suspectdevices.com:465 -starttls smtp
connect: Connection refused
connect:errno=111
```

Add ssl/tls to postfix for outgoing mail

```
root@naomi:/etc/postfix# nano master.cf
...
# =====
# service type private unpriv chroot wakeup maxproc command + args
# (yes) (yes) (no) (never) (100)
# =====
smtp inet n - y - - smtpd
#smtp inet n - y - 1 postscreen
#smtpd pass - - y - - smtpd
#dnsblog unix - - y - 0 dnsblog
#tlsproxy unix - - y - 0 tlsproxy
submission inet n - y - - smtpd
-o syslog_name=postfix/submission
-o smtpd_tls_security_level=encrypt
-o smtpd_sasl_auth_enable=yes
# -o smtpd_reject_unlisted_recipient=no
# -o smtpd_client_restrictions=$mua_client_restrictions
# -o smtpd_helo_restrictions=$mua_helo_restrictions
# -o smtpd_sender_restrictions=$mua_sender_restrictions
# -o smtpd_recipient_restrictions=
-o smtpd_relay_restrictions=permit_sasl_authenticated,reject
-o milter_macro_daemon_name=ORIGINATING
root@naomi:/etc/postfix# service postfix check
postfix/postfix-script: warning: group or other writable: /usr/lib/postfix/.sbin/lmt
postfix/postfix-script: warning: group or other writable: /usr/lib/postfix/.libpostfix-tls.so.1
postfix/postfix-script: warning: group or other writable: /usr/lib/postfix/.libpostfix-global.so.1
postfix/postfix-script: warning: group or other writable: /usr/lib/postfix/.libpostfix-master.so.1
postfix/postfix-script: warning: group or other writable: /usr/lib/postfix/.libpostfix-dns.so.1
postfix/postfix-script: warning: group or other writable: /usr/lib/postfix/.libpostfix-util.so.1
postfix/postfix-script: warning: group or other writable: /usr/lib/postfix/sbin/.lmt
root@naomi:/etc/postfix# service postfix reload
```

Link authentication to dovecot and enable auth server in dovecot. " apparently this can be avoided by installing a single package buried in ubuntu's documentation (g: Mail-Stack Delivery).

```
root@naomi:/etc/postfix# nano /etc/dovecot/conf.d/10-master.conf
...
#Postfix smtp-auth
unix_listener /var/spool/postfix/private/auth {
  mode = 0666
}

# Auth process is run as this user.
#user = $default_internal_user
}

service auth-worker {
  # Auth worker process is run as root by default, so that it can access
  # /etc/shadow. If this isn't necessary, the user should be changed to
  # $default_internal_user.
  user = root
}
...

root@naomi:/etc/postfix# nano main.cf
# TLS parameters
smtpd_tls_cert_file=/etc/ssl/certs/ssl-cert-snakeoil.pem
smtpd_tls_key_file=/etc/ssl/private/ssl-cert-snakeoil.key
smtpd_use_tls=yes
smtpd_tls_auth_only = yes
smtpd_sasl_type = dovecot
smtpd_sasl_path = private/auth
smtpd_sasl_auth_enable = yes
smtpd_recipient_restrictions = permit_sasl_authenticated permit_mynetworks reject_unauth_destination
```

Follow up on above errors

NOTE: the above errors are related to symlinks and not the files. Both debian and canonical aren't concerned about it and may or may not fix it at some point. <https://bugs.launchpad.net/ubuntu/+source/postfix/+bug/1728723>

eliminate pop3 as it isn't needed

```
mv /usr/share/dovecot/protocols.d/pop3d.protocol /usr/share/dovecot/pop3d.protocol.disabled
service dovecot reload
netstat -ta
```

5.103.2 SPF and openDKIM

Gmail currently requires that any email you send that isn't controlled by them use both SPF and DKIM.

What the hell is it?

According to linuxbabe <https://www.linuxbabe.com/mail-server/setting-up-dkim-and-spf>

SPF and DKIM are two types of TXT records in DNS that can help prevent email spoofing and ensure legitimate emails are delivered into the recipient's inbox instead of spam folder. If your domain is abused by email spoofing, then your emails are likely to landed in recipient's spam folder if they didn't add you in address book.

SPF (Sender Policy Framework) record specifies which hosts or IP addresses are allowed to send emails on behalf of a domain. You should allow only your own email server or your ISP's server to send emails for your domain.

_DKIM (DomainKeys Identified Mail) uses a private key to add a signature to emails sent from your domain. Receiving SMTP servers verify the signature by using the corresponding public key, which is published in your DNS manager. _

SPF

We only want to send email through a single server which is accomplished with the following record. Which needs to be added for each domain using the email server.

```
root@naomi:~# nano /etc/bind/zones/fromhell.hosts
... add the following ...
@ TXT "v=spf1 ip4:198.202.31.141 -all"
```

openDKIM

GOTCHAS

- convoluted and complex configuration involving 3 major services (dns, postfix, opendkim).
- postfix is chrooted and milter version is currently 6
- sample output from current opendkim-tools is wrong and requires manual correction.
- Relaying requires masquerading.

INSTALLATION

Install opendkim and edit configuration file

```
root@naomi:~# apt-get install opendkim opendkim-tools
root@naomi:~# nano /etc/opendkim.conf
... add/correct the following ...
Socket          local:/var/spool/postfix/var/run/opendkim/opendkim.sock
PidFile         /var/run/opendkim/opendkim.pid
Syslog          yes
UMask           002
UserID          opendkim
KeyTable        refile:/etc/opendkim/key.table
SigningTable    refile:/etc/opendkim/signing.table
ExternalIgnoreList refile:/etc/opendkim/trusted.hosts
InternalHosts   refile:/etc/opendkim/trusted.hosts
```

For each domain being handled create a signing key and add to dns zone files.

```
root@naomi:~# cd /etc/opendkim/keys/
root@naomi:/etc/opendkim/keys# opendkim-genkey -b 2048 -h rsa-sha256 -r -s 201807 -d suspectdevices.com -v
root@naomi:/etc/opendkim/keys# mv 201807.private suspectdevices.private
root@naomi:/etc/opendkim/keys# cat 201807.txt >>/etc/bind/zones/suspectdevices.hosts
```

Fix the error in dns entry and increment the zones serial number

```
root@naomi:/etc/opendkim/keys# nano /etc/bind/zones/suspectdevices.hosts
@           IN      SOA    dns1.digithink.com. don.digithink.com (
                2018072200 10800 3600 36000000 86400 )
...change.this. YYYYMMDDxx  ....
...
... and change h=rsa-sha256 to h=sha256 ...      ...as below...
201807._domainkey  IN      TXT      ( "v=DKIM1; h=sha256; k=rsa; s=email; "
                "p=MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA6ymvRLL+pEDThA6fMersYbr6dB5HKIFL4SM5F30RxcFmrYC//wm6/vrqWNft3AWy4zC7AQNiKyQGg7$
                "BUpxeL2b5GUhMrcZ+0heWwz7aF746IOY00IR4oMTFNP9a6hmmwBrLmna8ploFYUWCa2ETq/VYP6i14LU7P/yi8JhDMu4ZVI6ytLynBcLU42orcnWjwNLHqy/F3L$
```

Reload bind and check key

```
root@naomi:/etc/opendkim/keys# service bind9 reload
root@naomi:/etc/opendkim/keys# service bind9 status
● bind9.service - BIND Domain Name Server
   Loaded: loaded (/lib/systemd/system/bind9.service; enabled; vendor preset: enabled)
   Main PID: 127
   Status: active (running)
   Jul 25 22:35:15 naomi named[28512]: reloading zones succeeded
   ...
root@naomi:/etc/opendkim/keys# opendkim-testkey -d suspectdevices.com -s 201807 -vvv
opendkim-testkey: using default configfile /etc/opendkim.conf
opendkim-testkey: checking key '201807._domainkey.suspectdevices.com'
opendkim-testkey: key not secure .... ignore this ....
opendkim-testkey: key OK
```

Add entries to key.table signing.table and trusted hosts.

```
root@naomi:/etc/opendkim# nano key.table
fromhell      fromhell.com:201807:/etc/opendkim/keys/fromhell.private
suspectdevices suspectdevices.com:201807:/etc/opendkim/keys/suspectdevices.private
root@naomi:/etc/opendkim# nano signing.table
*@fromhell.com fromhell
*atfromhell.com fromhell
@suspectdevices.com suspectdevices
root@naomi:/etc/opendkim# nano trusted.hosts
127.0.0.1
::1
198.202.31.221
198.202.31.242
localhost
*.fromhell.com
*.suspectdevices.com
```

Configure socket file to communicate with postfix and add postfix to opendkim group.

```
root@naomi:~# mkdir -p /var/spool/postfix/var/run/opendkim
root@naomi:~# chown -R opendkim:opendkim /var/spool/postfix/var/run/opendkim
root@naomi:~# touch /var/spool/postfix/var/run/opendkim/opendkim.sock
root@naomi:~# chmod 775 /var/spool/postfix/var/run/opendkim/opendkim.sock
root@naomi:~# usermod -a -G opendkim postfix
root@naomi:~# nano /etc/default/opendkim
...
DAEMON_OPTS="-vvvv"
SOCKET="/local:/var/spool/postfix/var/run/opendkim/opendkim.sock"
RUNDIR=/var/spool/postfix/var/run/opendkim
USER=opendkim
GROUP=opendkim
PIDFILE=$RUNDIR/$NAME.pid
EXTRAFTER=
...
```

Add filter to postfix and restart both services.

```
root@naomi:~# nano /etc/postfix/main.cf
...
milter_protocol = 6
milter_default_action = accept
smtpd_milters = unix:/var/run/opendkim/opendkim.sock
non_smtpd_milters = unix:/var/run/opendkim/opendkim.sock
...
root@naomi:~# service opendkim reload
root@naomi:~# service postfix reload
```

Send test mail

```
root@naomi:~# echo "dkim test" |mail -testopendkim check-auth@verifier.port25.com
```

ADDING SIGNATURES TO RELAYED HOSTS

To relay mail from other hosts on the local networks requires the following additions to postfix's main.cf

```
root@naomi:~# nano /etc/postfix/main.cf
...
mynetworks = 127.0.0.0/8 [::ffff:127.0.0.0]/104 [::1]/128, 198.202.31.128/25
...
masquerade_domains = suspectdevices.com, fromhell.com
```

OPENDKIM/SPF LINKS

- <https://www.cioby.ro/2013/11/14/configuring-opendkim-to-sign-postfix-emails/>
- <https://linuxaria.com/howto/using-opendkim-to-sign-postfix-mails-on-debian>
- http://www.openspf.org/SPF_Record_Syntax
- <https://blog.whabash.com/posts/send-outbound-email-postfix-dkim-spf-ubuntu-16-04>
- <https://www.linode.com/docs/email/postfix/configure-spf-and-dkim-in-postfix-on-debian-8/>
- <https://www.linuxbabe.com/mail-server/setting-up-dkim-and-spf>
- <https://tools.ietf.org/html/rfc6376>
- <https://tweenpath.net/opendkim-postfix-smtp-relay-server-on-debian-7/>
- <https://qureshi.me/how-to-setup-postfixdkimspfmarc-on-ubuntu-plesk-onyx/>

5.103.3 Configure root/notification mail from other systems (esp bs2020)

Systems need to be able send email to notify us of issues such as security updates (apticron) etc. In order for email to be signed by opendkim and validated by spf the email needs to strip the hostname from mail sent from it before being relayed through the mail server.

```
root@bs2020:~# apt-get install mailutils apticron
... select satellite server when asked ...
root@bs2020:~# nano /etc/postfix/main.cf
... add the following ...
relayhost = naomi.suspectdevices.com
compatibility_level=2
masquerade_domains = suspectdevices.com
```

Since all systems will be striped of their machine names insure the full name of common accounts is made to be uniq

```
root@bs2020:~# chfn -f "Root at BS2020"
```

- http://www.postfix.org/STANDARD_CONFIGURATION_README.html
- <https://www.tecmint.com/setup-postfix-mail-server-smtp-using-null-client-on-centos/> *Todo:*
- I think postfix is a little heavy handed to run a null client. Investigate simpler secure solution.
- add amavis, and other filters linked in at <https://help.ubuntu.com/community/MailServer>
- make procmail do some work since its enabled by default
- make damned sure that it wont accept mail from the entire c-block

5.104 VideoRanchCloudServerConfiguration

Videoranch Cloud Server Configuration.

The purpose of this document is provide information on how gihon.orgs cloud server is currently configured and basic guidelines for maintaining it.

# Date	# Author	# Email	# Comments
28MAY16	Donald Delmar Davis	don@suspectdevices.com	Initial document

Background

We were asked to convert a 15 year old internet server running freebsd to the cloud. We started by setting up a staging server running Ubuntu 14.04 and migrating the users data and log files from the old server. This provided a backup of the original data and a place where we could work without having to pay for disk or bandwidth before deploying the final product. After a long process of porting all of the users and web sites that the server had served over the decades we began identifying which services, users, and domains were needed on the server. Given a much smaller set of users and web sites that were actually needed, we deployed an AWS image based on the AMI provided by the commercial entity which maintains Ubuntu. The active users users and web content have been installed on this server and the remainder has been archived to an external disk.

5.104.1 The Base Image

We chose to deploy an image provided by Canonical specifically for AWS "ubuntu-trusty-14.04-amd64-server-20150325 (ami-5189a661)" <http://cloud-images.ubuntu.com/releases/trusty/release-20150325/>

Adjustments to the image

The ubuntu user which provides a back door through which AWS allows users that it has authenticated to have root access to the instance. Unfortunately the ubuntu UID(1000) was already taken (jess) so it was moved to 999 and files owned by it were migrated as well.

```
chown --from=1000:1000 999:999 /. -Rv
```

Also the mail spool was somewhere new (/var/spool/mail) so I linked the new location back to /var/mail

Additions to the image

a lamp stack was added to the image using the "tasksel" package which bundles most services into supported configurations and deploys them along with all of their dependencies. (Note that the Ubuntu Cloud Image was already installed)

```
# tasksel
Package configuration
```

```

|-----| Software selection |-----|
| You can choose to install one or more of the following predefined collections of software. |
|
| Choose software to install:
```

```

|      [*] Basic Ubuntu server
|      [*] OpenSSH server
|      [ ] DNS server
|      [*] LAMP server
|      [*] Mail server
|      [*] PostgreSQL database
|      [ ] Print server
|      [ ] Samba file server
|      [ ] Tomcat Java server
|      [*] Ubuntu Cloud Image (instance)
|      [ ] Virtual Machine host
...
|
|      <Ok>
|

```

users and superusers

The following users were added to the system.

```

jess:x:1000:1000:Jessica Kent:/home/jess:/bin/csh
gepr:x:1053:1053:Glen E Ropella:/home/gepr:/bin/bash
don:x:1054:1054:Donald Delmar Davis:/home/don:/bin/bash
vic:x:1002:1002:Victoria Kennedy:/home/vic:/bin/bash
nez:x:1003:1003:Michael Nesmith:/home/nez:/bin/bash
vranch:x:1004:1004:Videoranch User:/home/vranch:/bin/bash
foreman:x:1005:1005:Videoranch Foreman:/home/foreman:/bin/tcsh
navajoslim:x:1007:1007:Navajo Slim:/home/navajoslim:/bin/bash
gihon:x:1017:1017:Gihon Foundation:/home/gihon:/bin/bash
vk:x:1021:1021:Victoria Kennedy:/home/vk:/bin/bash
vrresume:x:1024:1024:videoranch resume:/home/vrresume:/bin/bash
vak:x:1027:1027:victoria kennedy:/home/vak:/bin/tcsh
nezrays:x:1031:1031:nezrays:/usr/home/vranch/nezrays/www:/bin/sh
vr3d:x:1035:1035:VR3D:/home/vr3d:/bin/sh
staging:x:1041:1041:staging:/home/staging:/bin/bash
nesmith:x:1042:1042:nesmith:/home/nesmith:/bin/bash
director:x:1045:1045:Jessica Kent:/home/director:/bin/bash
petetest:x:1048:1048:petetest:/home/petetest:/bin/bash
mn:x:1022:1022:Michael Nesmith:/home/mn:/bin/bash

```

This had to be done manually as some of the original passwords were so old that their encryption methods were no longer supported. In cases where the users were less than a few years old the users passwords transferred to the new system seamlessly. In other cases the passwords will have to be reset by someone with root access.

```
ubuntu@cloud # passwd vranch
```

Their mail pools (/var/mail/), and home directories were copied over as well.

sudo privileges were enabled for members of the sudo group.

```

ubuntu@cloud # vigr
...
sudo:x:27:ubuntu,jess,foreman,don,gepr
...

```

5.104.2 Apache Configuration

In addition to the home directories of the remaining users the /home/vranch directory tree and /home/gihon were copied to the new server. The server configurations were ported to be as close to the originals as possible. (exceptions noted below)

The default server is set to www.gihon.com and is configured based on the original virtual-host. The php information and much about the apache server can be queried directly at <http://videoranch.com/test.php>

```

#ServerName www.gihon.com
<VirtualHost *:80>
    ServerName www.gihon.com
    ServerAlias gihon.com www.gihon.org gihon.org cloud.gihon.com
    ServerAdmin info@digitaloffspring.com
    DocumentRoot /home/gihon/www
    <Directory '/home/gihon/'>
        AllowOverride All
    </Directory>
    ScriptAlias /cgi-bin/ /home/gihon/cgi-bin
    CustomLog /home/gihon/logs/gihon-access_log common

```

```
ErrorLog /home/gihon/logs/gihon-error_log
</VirtualHost>
```

- Note that the log files are left in user space (off of /home) this allows clients to pull and view the log files in the same way that they update the content of their web site (ftp etc)
- Some configuration directives are no longer supported and are commented out.
- Extremely dangerous statements such as AllowOverides for the root directory were modified.

All other servers are named virtualhosts. The first of which is www.videoranch.com defined in /etc/apache2/sites-enabled/www.videoranch.com.conf

```
<VirtualHost *:80>
    ServerName www.videoranch.com
    ServerAlias videoranch.com www.videoranch.com
#   Header append p3p 'CP="OTI DSP COR CUR UNI" polyref="/w3c/p3policy.xml"'
    ServerAdmin info@digitaloffspring.com
    DocumentRoot /home/vranch/videoranch/www
    ScriptAlias /cgi-bin/ /home/vranch/videoranch3d/cgi-bin/
    ErrorLog /home/vranch/logs/www.videoranch.com-error_log
    CustomLog /home/vranch/logs/www.videoranch.com-access_log common
    <Directory /home/vranch/videoranch/www>
        Options Indexes FollowSymLinks
        AllowOverride All
    </Directory>
</VirtualHost>
```

5.104.3 Pro-ftp Configuration

We configured proftpd (which we vetted as a viable and secure ftp daemon) as closely as possible to the original configuration on the old server. Because AWS instances are in their own private network and access has to be explicitly allowed you must specify the PASV ports in /etc/proftpd/proftpd.conf. These ports must be opened up in the "Security Group" configuration as well.

```
# In some cases you have to specify passive ports range to by-pass
# firewall limitations. Ephemeral ports can be used for that, but
# feel free to use a more narrow range.
PassivePorts          49152 49153
```

Ftp in its native form is insecure and so we would prefer to have configured an SSL certificate and require TLS for all ftp requests. We were able to verify that SFTP (ftp provided by ssh).

5.104.4 Network and "Security Group" configuration

The AWS instance is placed in a private network. This network provides the instance a private ip through dhcp. For this reason the main interface is configured as follows in /etc/networks/interfaces.d/eth0

```
# The primary network interface
auto eth0
iface eth0 inet dhcp
```

This address is attached to the outside world via an "Elastic" ip (52.34.143.142). To connect the external traffic to the private address you have to create a "Security group" and define the rules which allow traffic in and out of the private network.

- INBOUND RULES * || [# protocol|# family|# port|# allow from| |-----|-----|-----| | HTTP | TCP | 80 | 0.0.0.0/0 | | SSH | TCP | 22 | 0.0.0.0/0 | | SMTP | TCP | 25 | 0.0.0.0/0 | | Custom TCP Rule | TCP | 20 - 21 | 0.0.0.0/0 | | IMAP | TCP | 143 | 0.0.0.0/0 | | Custom TCP Rule | TCP | 49152 - 49153 | 0.0.0.0/0 | | HTTPS | TCP | 443 | 0.0.0.0/0 |

Outbound rules allow all outgoing traffic.

5.104.5 Unused Capabilities

MySQL and PostgreSQL

While the M in LAMP is MySQL, Many developers prefer Postgres which is much more standards oriented and robust. Both databases are available and PHP is configured for them. At one point mysql was on the old server however neither gihon nor the model files served by videoranch.com seemed to use it. _ Note that if either database is used a mechanism to back up the data must also be implimented_

Postfix and Dovecot

The standard SMTP (email) server for most current operating systems is Postfix. The Mail server task also includes Dovecot which provides both POP and IMAP servers for clients to download any mail still on the server. To use the pop server will require the addition of the ports for pop (110) to be added to the security group configuration. _These servers are not currently configured. _

5.104.6 Log Rotation Configuration

On the previous server most log files were larger than the content being provided. Ubuntu provides a log rotation utility designed to compress and delete logs in a reasonable manner preventing them from consuming system resources over time. Since the apache logs on this system are in "user space" and not under /var/log/apache2 their location needed to be configured.

Here is the section added to /etc/logrotate.d/apache2 for the gihon.com

```
/home/gihon/logs/*_log {
    weekly
    missingok
    rotate 52
    compress
    delaycompress
    notifempty
    create 640 root adm
    sharedscripts
    postrotate
        if /etc/init.d/apache2 status > /dev/null ; then \
            /etc/init.d/apache2 reload > /dev/null; \
        fi;
    endscript
    prerotate
        if [ -d /etc/logrotate.d/httpd-prerotate ]; then \
            run-parts /etc/logrotate.d/httpd-prerotate; \
        fi; \
    endscript
}
```

5.104.7 unattended upgrades (security only)

The system is configured to automatically install security upgrades as released by the operating system. *In the event that an error occurs mail is sent to the foreman account.*

5.104.8 Operations Guide

Given the state of the previous system the soundest approach is to automate as much of the systems upkeep as possible. Log rotation and unattended system upgrades along with other minor adjustments (turning on apt's auto-remove for instance) should enable us to think of the box more as an appliance.

Backing up Server work with Live Snapshots

AWS allows a server to be backed up while running. These snapshots can be run up as separate servers (for development or to do a major release upgrade) Or they can be reattached to an existing instance (in the case of disaster or compromise). Please make a snapshot of the server whenever significant work has been done to it.

Backing up your data

Since the servers web content is in the user space. Log files, websites and other data served should be copied to a local server preferably one behind a firewall. *In particular Gihon should take care to keep updated copies of /home/gihon and /home/vranch*

Accessing the server

Privileged access can be granted through AWS to the Ubuntu user. For instructions on how to do this see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AccessingInstancesLinux.html>. The server has been configured to allow ssh access directly.

```
$ ssh www.videoranch.com
Welcome to Ubuntu 14.04.4 LTS (GNU/Linux 3.13.0-85-generic x86_64)
```

```
* Documentation: https://help.ubuntu.com/

System information as of Mon Apr 11 18:12:10 UTC 2016

System load:  0.0          Processes:      139
Usage of /:   69.8% of 29.39GB  Users logged in:  1
Memory usage: 28%          IP address for eth0: 172.31.16.108
Swap usage:   0%

Graph this data and manage this system at:
https://landscape.canonical.com/

Get cloud support with Ubuntu Advantage Cloud Guest:
http://www.ubuntu.com/business/services/cloud

0 packages can be updated.
0 updates are security updates.

You have new mail.
Last login: Mon Apr 11 16:11:57 2016 from 71-34-91-188.ptld.qwest.net
don@cloud:~$
```

References

- why ubuntu? <https://insights.ubuntu.com/2014/04/15/ubuntu-14-04-lts-the-cloud-platform-of-choice/>
- <https://www.digitalocean.com/community/tutorials/how-to-configure-logging-and-log-rotation-in-apache-on-an-ubuntu-vps>
- <https://help.ubuntu.com/lts/serverguide/automatic-updates.html>
- <https://anturis.com/linux-server-maintenance-checklist/>

5.105 fixing

trac database The database created when upgrading the trac site does not work with wikiprintbook.

Here is a workaround while I debug the issue.

```
root@herbert:~# mkdir /tmp/docsdump
root@herbert:~# trac-admin /var/trac/serverdocs/env wiki dump /tmp/docsdump/
WikiNewPage => /tmp/docsdump/WikiNewPage
PlatformIO => /tmp/docsdump/PlatformIO
GoodByeOpenstack => /tmp/docsdump/GoodByeOpenstack
Esp8266 => /tmp/docsdump/Esp8266
Mullein => /tmp/docsdump/Mullein
PageTemplates => /tmp/docsdump/PageTemplates
DockerInstallNotes => /tmp/docsdump/DockerInstallNotes
Ubuntu18.04Notes => /tmp/docsdump/Ubuntu18.04Notes
DL380RaidController => /tmp/docsdump/DL380RaidController
ResliverFailureNotes => /tmp/docsdump/ResliverFailureNotes
InitialImpressions => /tmp/docsdump/InitialImpressions
InterMapTxt => /tmp/docsdump/InterMapTxt
RecentChanges => /tmp/docsdump/RecentChanges
OpenWRTonLinkSysEA3500 => /tmp/docsdump/OpenWRTonLinkSysEA3500
OpenWRT => /tmp/docsdump/OpenWRT
ZFSDiskReplacement => /tmp/docsdump/ZFSDiskReplacement
InterWiki => /tmp/docsdump/InterWiki
BleedingEdgeServer => /tmp/docsdump/BleedingEdgeServer
DiskRecovery => /tmp/docsdump/DiskRecovery
CamelCase => /tmp/docsdump/CamelCase
WikiStart => /tmp/docsdump/WikiStart
7900NWashburne => /tmp/docsdump/7900NWashburne
Feurig => /tmp/docsdump/Feurig
ILO3Notes => /tmp/docsdump/ILO3Notes
SandBox => /tmp/docsdump/SandBox
LEDE => /tmp/docsdump/LEDE
FunWithLinuxDisks => /tmp/docsdump/FunWithLinuxDisks
OpenVPNOnLEDE => /tmp/docsdump/OpenVPNOnLEDE
ZFSPMirroredFromExisting => /tmp/docsdump/ZFSPMirroredFromExisting
ZFSNightmaresPorted2Linux => /tmp/docsdump/ZFSNightmaresPorted2Linux
LXDContainerWithDockerNotes => /tmp/docsdump/LXDContainerWithDockerNotes
MigratingServicesToLXC => /tmp/docsdump/MigratingServicesToLXC
kb2018InstallBashHistory => /tmp/docsdump/kb2018InstallBashHistory
MigrateUsers => /tmp/docsdump/MigrateUsers
ZFSHotSwappingMirrorsOnLivePools => /tmp/docsdump/ZFSHotSwappingMirrorsOnLivePools
CloudServerConfiguration => /tmp/docsdump/CloudServerConfiguration
ContainerShipInstallation => /tmp/docsdump/ContainerShipInstallation
OpenWrtE900FirmwareBuild => /tmp/docsdump/OpenWrtE900FirmwareBuild
InterTrac => /tmp/docsdump/InterTrac
BS2020InstallNotes => /tmp/docsdump/BS2020InstallNotes
TitleIndex => /tmp/docsdump/TitleIndex
UbuntuMailServerSetup => /tmp/docsdump/UbuntuMailServerSetup
Annie => /tmp/docsdump/Annie
GlassesOkay => /tmp/docsdump/GlassesOkay
CloudServerDocs => /tmp/docsdump/CloudServerDocs
TicketQuery => /tmp/docsdump/TicketQuery
OperationsGuide => /tmp/docsdump/OperationsGuide
DiskLayoutOnBS2020 => /tmp/docsdump/DiskLayoutOnBS2020
LXDContainersWithProfile => /tmp/docsdump/LXDContainersWithProfile
Nigel => /tmp/docsdump/Nigel
Idrac6 => /tmp/docsdump/Idrac6
AutoMatingContainerUpdates => /tmp/docsdump/AutoMatingContainerUpdates
OpenWRTonMR3020 => /tmp/docsdump/OpenWRTonMR3020
NewTracContainer => /tmp/docsdump/NewTracContainer
SuspectDevices => /tmp/docsdump/SuspectDevices
```

```
SystemUpdates => /tmp/docsdump/SystemUpdates
CaptiveRaidController => /tmp/docsdump/CaptiveRaidController
```

- clear out any existing pages on the disposable wiki site

```
root@herbert:~# trac-admin /var/trac/devel/env wiki remove *
```

5.105.1 Deleted pages

RecentChanges InterWiki TicketQuery CamelCase WikiStart PageTemplates InterTrac SandBox TitleIndex InterMapTxt OperationsGuide

- load the pages onto the new site. (may be missing a step for the images)

```
root@herbert:~# trac-admin /var/trac/devel/env wiki load /tmp/.ICE-unix/ env/ .Test-unix/ files/ .X11-unix/ netplan_141i3qzp/
.XIM-unix/ systemd-private-a1dddc0dcb0479fad96fa3c064e61e2-apache2.service-Icqav1/ .font-unix/ systemd-private-
a1dddc0dcb0479fad96fa3c064e61e2-systemd-resolved.service-gIdUDr/ docsdump/ trackage28nov18.tgz
root@herbert:~# trac-admin /var/trac/devel/env wiki load /tmp/docsdump/ ZFSMirroredFromExisting imported from /tmp/
docsdump/ZFSMirroredFromExisting BS2020InstallNotes imported from /tmp/docsdump/BS2020InstallNotes SandBox
imported from /tmp/docsdump/SandBox MigrateUsers imported from /tmp/docsdump/MigrateUsers UbuntuMailServerSetup
imported from /tmp/docsdump/UbuntuMailServerSetup Idrac6 imported from /tmp/docsdump/Idrac6 DockerInstallNotes
imported from /tmp/docsdump/DockerInstallNotes WikiNewPage imported from /tmp/docsdump/WikiNewPage WikiStart
imported from /tmp/docsdump/WikiStart kb2018InstallBashHistory imported from /tmp/docsdump/kb2018InstallBashHistory
Feurig imported from /tmp/docsdump/Feurig PageTemplates imported from /tmp/docsdump/PageTemplates
ZFSHotSwappingMirrorsOnLivePools imported from /tmp/docsdump/ZFSHotSwappingMirrorsOnLivePools ILO3Notes
imported from /tmp/docsdump/ILO3Notes SystemUpdates imported from /tmp/docsdump/SystemUpdates OperationsGuide
imported from /tmp/docsdump/OperationsGuide CaptiveRaidController imported from /tmp/docsdump/CaptiveRaidController
DL380RaidController imported from /tmp/docsdump/DL380RaidController OpenWRTonMR3020 imported from /tmp/
docsdump/OpenWRTonMR3020 OpenWRT imported from /tmp/docsdump/OpenWRT RecentChanges imported from /tmp/
docsdump/RecentChanges LEDE imported from /tmp/docsdump/LEDE CloudServerConfiguration imported from /tmp/
docsdump/CloudServerConfiguration GlassesOkay imported from /tmp/docsdump/GlassesOkay OpenWRTonLinkSysEA3500
imported from /tmp/docsdump/OpenWRTonLinkSysEA3500 AutoMatingContainerUpdates imported from /tmp/docsdump/
AutoMatingContainerUpdates DiskLayoutOnBS2020 imported from /tmp/docsdump/DiskLayoutOnBS2020 CamelCase
imported from /tmp/docsdump/CamelCase MigratingServicesToLXC imported from /tmp/docsdump/MigratingServicesToLXC
SuspectDevices imported from /tmp/docsdump/SuspectDevices Esp8266 imported from /tmp/docsdump/Esp8266
CloudServerDocs imported from /tmp/docsdump/CloudServerDocs Annie imported from /tmp/docsdump/Annie
GoodByeOpenstack imported from /tmp/docsdump/GoodByeOpenstack TicketQuery imported from /tmp/docsdump/
TicketQuery OpenWrtE900FirmwareBuild imported from /tmp/docsdump/OpenWrtE900FirmwareBuild FunWithLinuxDisks
imported from /tmp/docsdump/FunWithLinuxDisks InterMapTxt imported from /tmp/docsdump/InterMapTxt
Ubuntu18.04Notes imported from /tmp/docsdump/Ubuntu18.04Notes ZFSDiskReplacement imported from /tmp/docsdump/
ZFSDiskReplacement DiskRecovery imported from /tmp/docsdump/DiskRecovery InterTrac imported from /tmp/docsdump/
InterTrac NewTracContainer imported from /tmp/docsdump/NewTracContainer ZFSNightmaresPorted2Linux imported from
/tmp/docsdump/ZFSNightmaresPorted2Linux ContainerShipInstallation imported from /tmp/docsdump/
ContainerShipInstallation PlatformIO imported from /tmp/docsdump/PlatformIO Nigel imported from /tmp/docsdump/Nigel
TitleIndex imported from /tmp/docsdump/TitleIndex LXDCContainerWithDockerNotes imported from /tmp/docsdump/
LXDCContainerWithDockerNotes BleedingEdgeServer imported from /tmp/docsdump/BleedingEdgeServer OpenVPNOnLEDE
imported from /tmp/docsdump/OpenVPNOnLEDE InterWiki imported from /tmp/docsdump/InterWiki
LXDCContainersWithProfile imported from /tmp/docsdump/LXDCContainersWithProfile Mullein imported from /tmp/docsdump/
Mullein ResilverFailureNotes imported from /tmp/docsdump/ResilverFailureNotes InitialImpressions imported from /tmp/
docsdump/InitialImpressions 7900NWashburne imported from /tmp/docsdump/7900NWashburne root@herbert:~#
```

Then go to the [/devel devel] site and print the book.

5.106 ZFS Disk Replacement

The process of replacing mirrored zfs disks is fairly simple. The changes are done by zpool attach and detach.

```
zpool detach <pool> <disk-id>
zpool attach <pool> <disk-id-to-mirror> <disk-id-mirrored-to>
```

The heavy lifting is done by zfs itself.

5.106.1 process

PREP

- use the [pdf article link](#) to print this before going down
- If possible pre wipe and check the disks on a separate linux machine (*note: /dev/sdf is an placeholder for the disk mounted on that system*)

```
``` root@homebox:~# wipefs -af --backup /dev/sdf /dev/sdf: 8 bytes were erased at offset 0x00000200 (gpt): 45 46 49 20 50 41
52 54 /dev/sdf: 8 bytes were erased at offset 0x222ee64e00 (gpt): 45 46 49 20 50 41 52 54 /dev/sdf: 2 bytes were erased at offset
0x000001fe (PMBR): 55 aa /dev/sdc: calling ioctl to re-read partition table: Success root@homebox:~# fdisk /dev/sdf Command
(m for help): g
```

```
Created a new GPT disklabel (GUID: EBC5A0C9-E871-544F-A8EA-E31FCA655F9C).
```

```
Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
```

```
root@homebox:~# badblocks /dev/sdf
....
```

```
* insure that you can ssh into the box
```

```
On site
```

```
The following assumes you have escalated to root privileges (sudo bash), in this case we are replacing /dev/sdc and /dev/sdd in the pool named 'level'
```

```
* check for the correct disk.
```

```
The following should cause the disk to light up\
(CTRL> C when you have identified the disk. Careful with the if/of here).
...
root@bs2020:~# dd if=/dev/sdc of=/dev/null
```

- find the disk in the pool.

```
``` root@bs2020:~# zpool status pool: devel state: ONLINE scan: resilvered 9.95G in 0h4m with 0 errors on Sat Nov 10 22:00:41
2018 config:
```

```
NAME                STATE      READ WRITE CKSUM
devel                ONLINE    0    0    0
  mirror-0           ONLINE    0    0    0
    scsi-35000c50054fee503 ONLINE    0    0    0
    scsi-35000c5005501b45b ONLINE    0    0    0

errors: No known data errors

...
root@bs2020:~# ls -ls /dev/disk/by-id|grep scsi|grep -v "\-part"
0 lrwxrwxrwx 1 root root 9 Nov 10 21:22 scsi-350000395a8336d34 -> ../../sde
0 lrwxrwxrwx 1 root root 9 Nov 10 21:22 scsi-35000c50054fee503 -> ../../sdd
0 lrwxrwxrwx 1 root root 9 Nov 10 21:56 scsi-35000c5005501b45b -> ../../sdc
0 lrwxrwxrwx 1 root root 9 Nov 10 21:22 scsi-35000cca00b33a264 -> ../../sdf
0 lrwxrwxrwx 1 root root 9 Nov 10 21:22 scsi-3600508e00000000069cf3977618f1408 -> ../../sdg
root@bs2020:~#
```

```
_We notice above that the disk we are looking for is scsi-35000c5005501b45b_
```

```
* detach the disk from the pool.
```

```
...
root@bs2020:~# zpool detach devel scsi-35000c5005501b45b
root@bs2020:~# zpool status
```

```

pool: devel
state: ONLINE
scan: resilvered 9.95G in 0h4m with 0 errors on Sat Nov 10 22:00:41 2018
config:

    NAME                STATE      READ WRITE CKSUM
    devel                ONLINE    0    0    0
    scsi-35000c50054fee503  ONLINE    0    0    0

errors: No known data errors

pool: infra
...
root@bs2020:~#

```

- even if expanding the disk size insure that autoexpand is off.

```
root@bs2020:~# zpool set autoexpand=off devel
```

```

* Swap out the old disk with the new one.

* find the new disk's id.

...
root@bs2020:~# partprobe
root@bs2020:~# ls -ls /dev/disk/by-id/|grep sdc
0 lrwxrwxrwx 1 root root  9 Nov 10 21:56 scsi-xxxxxxxxxxxxxxxx -> ../../sdc
...
0 lrwxrwxrwx 1 root root  9 Nov 10 21:56 xxx-xxxxxxxxxxxxxxxx -> ../../sdc

```

- *If the drive id does not change reboot the server*
- attach the new disk to the zfs pool (*scsi-xxxxxxxxxxxxxxxx is the new id from the above step*)

```
root@bs2020:~# zpool attach devel scsi-35000c50054fee503 scsi-xxxxxxxxxxxxxxxx
```

```

* wait for pool to resilver

...
root@bs2020:~# zpool status
pool: devel
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Sat Nov 10 21:56:04 2018
      8.54G scanned out of 9.95G at 35.5M/s, 0h0m to go
      8.54G resilvered, 85.85% done
config:

    NAME                STATE      READ WRITE CKSUM
    devel                ONLINE    0    0    0
    mirror-0             ONLINE    0    0    0
    scsi-35000c50054fee503  ONLINE    0    0    0
    scsi-35000c5005501b45b  ONLINE    0    0    0 (resilvering)

errors: No known data errors

pool: infra
...

root@bs2020:~# zpool status
.... repeat until finished resilvering ....
root@bs2020:~# zpool status
pool: devel
state: ONLINE
scan: scrub repaired 0B in 0h4m with 0 errors on Sat Nov 10 21:58:04 2018
config:

    NAME                STATE      READ WRITE CKSUM
    devel                ONLINE    0    0    0
    mirror-0             ONLINE    0    0    0
    scsi-35000cca00b33a264  ONLINE    0    0    0
    scsi-350000395a8336d34  ONLINE    0    0    0

errors: No known data errors
...

```

- if expanding disk check for new size and if not expand it

...

zfs list and check for larger disk pool

* repeat process for disk in bay below (we already know its old id from above).

...

```
root@bs2020:~# dd if=/dev/sdd of=/dev/null
root@bs2020:~# zpool detach devel scsi-35000c50054fee503
... swap disks ...
root@bs2020:~# partprobe
root@bs2020:~# ls -ls /dev/disk/by-id|grep sdd
0 lrwxrwxrwx 1 root root 9 Nov 10 21:56 scsi-yyy-yyyy-yyyy-yyyy-> ../../sdd
... reboot if necessary ...
root@bs2020:~# wipefs -a /dev/sdd
...
root@bs2020:~# fdisk /dev/sdd
...
root@bs2020:~# zpool attach devel scsi-xxxxxxxxxxxxx scsi-yyy-yyyy-yyyy-yyyy
... wait for resilver...
```

- use the process below to grow disks to new size

```
# zpool set autoexpand=on devel
# zpool online -e devel scsi-xxxxxxxxxxxxxxxxxxxxx
# zpool online -e devel scsi-yyy-yyyy-yyyy-yyyy-yyyy
# zpool set autoexpand=off devel
```

references

- <https://tomasz.korwel.net/2014/01/03/growing-zfs-pool/>
- <https://jsosic.wordpress.com/2013/01/01/expanding-zfs-zpool-raid/>
- <https://serverfault.com/questions/5336/how-do-i-make-linux-recognize-a-new-sata-dev-sda-drive-i-hot-swapped-in-without>

5.107 HOLY FUCKING AWESOME!!!!

Watch while I add a fresh disk as a mirror, resilver the pool and remove and repartition the original disk while the container using the pool is still running!!! ... make this into a structured document ...

```

root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4dev      ONLINE       0     0     0
      sdd1       ONLINE       0     0     0
      sdf        ONLINE       0     0     0
      sde        ONLINE       0     0     0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: scrub repaired 0 in 0h2m with 0 errors on Sun Aug 12 00:27:02 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4infra    ONLINE       0     0     0
      sdal       ONLINE       0     0     0

errors: No known data errors
root@bs2020:~# zpool add -n lxd4infra mirror sdb
invalid vdev specification: mirror requires at least 2 devices
root@bs2020:~# zpool add -n lxd4infra mirror sdal sdb
invalid vdev specification
use '-f' to override the following errors:
/dev/sdal is part of active pool 'lxd4infra'
/dev/sdb does not contain an EFI label but it may contain partition
information in the MBR.
root@bs2020:~# mklabel GPT /dev/sdb
bash: mklabel: command not found
root@bs2020:~# parted /dev/sdb
bash: parted: command not found
root@bs2020:~# gparted /dev/sdb
bash: gparted: command not found
root@bs2020:~# zpool add -nf lxd4infra mirror sdal sdb
invalid vdev specification
the following errors must be manually repaired:
/dev/sdal is part of active pool 'lxd4infra'
root@bs2020:~# zpool add -nf lxd4infra mirror sdb sdal
invalid vdev specification
the following errors must be manually repaired:
/dev/sdal is part of active pool 'lxd4infra'
root@bs2020:~# zpool add -nf lxd4infra mirror sdb
invalid vdev specification: mirror requires at least 2 devices
root@bs2020:~# zpool add -nf lxd4infra sdal mirror sdb
invalid vdev specification: mirror requires at least 2 devices
root@bs2020:~# zpool attach -n sdal sdb
invalid option 'n'
usage:
  attach [-f] [-o property=value] <pool> <device> <new-device>
root@bs2020:~# zpool attach sdal sdb
missing <new_device> specification
usage:
  attach [-f] [-o property=value] <pool> <device> <new-device>
root@bs2020:~# zpool attach lxd4infra sdal sdb
invalid vdev specification
use '-f' to override the following errors:
/dev/sdb does not contain an EFI label but it may contain partition
information in the MBR.
root@bs2020:~# gparted
bash: gparted: command not found
root@bs2020:~# parted
bash: parted: command not found
root@bs2020:~# apt-get install parted
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  libparted2
Suggested packages:
  libparted-dev libparted-il18n parted-doc
The following NEW packages will be installed:
  libparted2 parted
0 upgraded, 2 newly installed, 0 to remove and 0 not upgraded.
Need to get 158 kB of archives.
After this operation, 520 kB of additional disk space will be used.
Do you want to continue? [Y/n]
Get:1 http://us.archive.ubuntu.com/ubuntu xenial-updates/main amd64 libparted2 amd64 3.2-15ubuntu0.1 [115 kB]
Get:2 http://us.archive.ubuntu.com/ubuntu xenial-updates/main amd64 parted amd64 3.2-15ubuntu0.1 [42.4 kB]

```



```

Fetched 158 kB in 0s (277 kB/s)
Selecting previously unselected package libparted2:amd64.
(Reading database ... 32152 files and directories currently installed.)
Preparing to unpack .../libparted2_3.2-15ubuntu0.1_amd64.deb ...
Unpacking libparted2:amd64 (3.2-15ubuntu0.1) ...
Selecting previously unselected package parted.
Preparing to unpack .../parted_3.2-15ubuntu0.1_amd64.deb ...
Unpacking parted (3.2-15ubuntu0.1) ...
Processing triggers for libc-bin (2.23-0ubuntu10) ...
Processing triggers for man-db (2.7.5-1) ...
Setting up libparted2:amd64 (3.2-15ubuntu0.1) ...
Setting up parted (3.2-15ubuntu0.1) ...
Processing triggers for libc-bin (2.23-0ubuntu10) ...
root@bs2020:~# parted /dev/sdb
GNU Parted 3.2
Using /dev/sdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mklabel GPT
(parted) w
      align-check TYPE N           check partition N for TYPE(min|opt) alignment
      help [COMMAND]              print general help, or help on COMMAND
      mklabel,mktable LABEL-TYPE  create a new diskLabel (partition table)
      mkpart PART-TYPE [FS-TYPE]  make a partition
      name NUMBER NAME            name partition NUMBER as NAME
      print [devices|free|list,all|NUMBER] display the partition table, available devices, free space, all found partitions, or a particular partition
      quit                        exit program
      rescue START END            rescue a lost partition near START and END
      resizepart NUMBER END        resize partition NUMBER
      rm NUMBER                    delete partition NUMBER
      select DEVICE                choose the device to edit
      disk_set FLAG STATE          change the FLAG on selected device
      disk_toggle [FLAG]           toggle the state of FLAG on selected device
      set NUMBER FLAG STATE        change the FLAG on partition NUMBER
      toggle [NUMBER [FLAG]]       toggle the state of FLAG on partition NUMBER
      unit UNIT                    set the default unit to UNIT
      version                      display the version number and copyright information of GNU Parted
(parted) q
Information: You may need to update /etc/fstab.

root@bs2020:~# zpool attach lxd4infra sda1 sdb
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

      NAME      STATE    READ WRITE CKSUM
      lxd4dev    ONLINE    0   0   0
        sdd1     ONLINE    0   0   0
        sdf      ONLINE    0   0   0
        sde      ONLINE    0   0   0

errors: No known data errors

pool: lxd4infra
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
       continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Tue Sep  4 09:05:14 2018
      182M scanned out of 5.38G at 10.7M/s, 0h8m to go
      181M resilvered, 3.30% done
config:

      NAME      STATE    READ WRITE CKSUM
      lxd4infra  ONLINE    0   0   0
        mirror-0  ONLINE    0   0   0
          sda1    ONLINE    0   0   0
          sdb     ONLINE    0   0   0 (resilvering)

errors: No known data errors
root@bs2020:~# packet_write_wait: Connection to 198.202.31.242: Broken pipe
steve:~ don$ ssh feurig@bs2020.suspectdevices.com
Welcome to Ubuntu 16.04.5 LTS (GNU/Linux 4.4.0-96-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

0 packages can be updated.
0 updates are security updates.

New release '18.04.1 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

Last login: Tue Sep  4 08:26:28 2018 from 75.164.203.77
feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# packet_write_wait: Connection to 198.202.31.242: Broken pipe
steve:~ don$ ssh feurig@bs2020.suspectdevices.com
Welcome to Ubuntu 16.04.5 LTS (GNU/Linux 4.4.0-96-generic x86_64)

```

```

* Documentation:  https://help.ubuntu.com
* Management:    https://landscape.canonical.com
* Support:       https://ubuntu.com/advantage

0 packages can be updated.
0 updates are security updates.

New release '18.04.1 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

Last login: Wed Sep  5 16:10:53 2018 from 75.164.203.77
feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# packet_write_wait: Connection to 198.202.31.242: Broken pipe
steve:~ don$
steve:~ don$ ssh feurig@bs2020.suspectdevices.com
Welcome to Ubuntu 16.04.5 LTS (GNU/Linux 4.4.0-96-generic x86_64)

* Documentation:  https://help.ubuntu.com
* Management:    https://landscape.canonical.com
* Support:       https://ubuntu.com/advantage

0 packages can be updated.
0 updates are security updates.

New release '18.04.1 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

Last login: Wed Sep  5 18:56:14 2018 from 75.164.203.77
feurig@bs2020:~$ sudo bash
[sudo] password for feurig:
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

    NAME      STATE    READ WRITE CKSUM
    lxd4dev    ONLINE   0   0   0
      sdd1     ONLINE   0   0   0
      sdf      ONLINE   0   0   0
      sde      ONLINE   0   0   0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: resilvered 5.38G in 0h6m with 0 errors on Tue Sep  4 09:11:31 2018
config:

    NAME      STATE    READ WRITE CKSUM
    lxd4infra  ONLINE   0   0   0
      mirror-0  ONLINE   0   0   0
        sdal    ONLINE   0   0   0
        sdb     ONLINE   0   0   0

errors: No known data errors
root@bs2020:~# zpool detach -n lxd4infra sdal
invalid option 'n'
usage:
    detach <pool> <device>
root@bs2020:~# zpool detach lxd4infra sdal
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:

    NAME      STATE    READ WRITE CKSUM
    lxd4dev    ONLINE   0   0   0
      sdd1     ONLINE   0   0   0
      sdf      ONLINE   0   0   0
      sde      ONLINE   0   0   0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: resilvered 5.38G in 0h6m with 0 errors on Tue Sep  4 09:11:31 2018
config:

    NAME      STATE    READ WRITE CKSUM
    lxd4infra  ONLINE   0   0   0
      sdb     ONLINE   0   0   0

errors: No known data errors
root@bs2020:~# gparted /dev/sda
bash: gparted: command not found
root@bs2020:~# parted /dev/sda
GNU Parted 3.2
Using /dev/sda
Welcome to GNU Parted! Type 'help' to view a list of commands.

```

```
(parted) mklabel GPT
Warning: The existing disk label on /dev/sda will be destroyed and all data on this disk will be lost. Do you want to continue?
Yes/No? y
(parted) q
Information: You may need to update /etc/fstab.
```

```
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4dev	ONLINE	0	0	0
sdd1	ONLINE	0	0	0
sdf	ONLINE	0	0	0
sde	ONLINE	0	0	0

```
errors: No known data errors
```

```
pool: lxd4infra
state: ONLINE
scan: resilvered 5.38G in 0h6m with 0 errors on Tue Sep 4 09:11:31 2018
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4infra	ONLINE	0	0	0
sdb	ONLINE	0	0	0

```
errors: No known data errors
```

```
root@bs2020:~# zpool attach lxd4infra sdb sda
```

```
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h8m with 0 errors on Sun Aug 12 00:32:48 2018
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4dev	ONLINE	0	0	0
sdd1	ONLINE	0	0	0
sdf	ONLINE	0	0	0
sde	ONLINE	0	0	0

```
errors: No known data errors
```

```
pool: lxd4infra
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Thu Sep 6 09:24:09 2018
69.8M scanned out of 5.42G at 5.37M/s, 0h17m to go
67.9M resilvered, 1.26% done
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4infra	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
sdb	ONLINE	0	0	0
sda	ONLINE	0	0	0 (resilvering)

```
errors: No known data errors
```

```
root@bs2020:~#
```

5.108 ZFS Mirrored data on existing file server

5.108.1 Adding zfs mirror to existing data

On Annie, the Home File Server we have a pair of matched 2T sata disks, one of which contains the majority of the shared data. We want to convert these to a mirrored disk using ZFS (thereby securing the existing data). Rather than using entire disks the disks should be partitioned so that they are bootable and can contain a fresh os installation.

note: the following assumes we have installed some prerequisites...

```
root@annie:~# apt-get install zfsutils-linux parted nfs-kernel-server zfs-initramfs
```

First we wipe and partition the unused disk.

```
root@annie:~# df -k
Filesystem      1K-blocks      Used Available Use% Mounted on
...
/dev/sdc1        1922728820 905645512 919391248  50% /export
/dev/sdd1        1921802520   77852 1824032608    1% /archive
...
root@annie:~# umount /archive
... adjust /etc/fstab if necessary ...
root@annie:~# parted /dev/sdd1
(parted) mklabel gpt
Warning: Partition(s) on /dev/sdd are being used.
Ignore/Cancel? I
(parted) mkpart zfs zfs 0% -100512MB
(parted) mkpart efi fat32 -100512MB -100000MB
(parted) mkpart lnx ext2 -100000MB 100%
(parted) set 2 boot on
root@annie:~# reboot
```

Create a zfs pool on the first partition

```
root@annie:~# zpool create basement -f /dev/disk/by-id/wwn-0x5000039ff3c899c1-part1
root@annie:~# zpool list
NAME      SIZE  ALLOC  FREE  EXPANDSZ  FRAG    CAP  DEDUP  HEALTH  ALTROOT
basement  1.72T   865G   895G        -         0%   49%  1.00x  ONLINE  -
root@annie:~# df -k
...
/dev/sdc1        1922728820 905645512 919391248  50% /export
...
basement         1787821824      128 1787821696    1% /basement
root@annie:~# mkdir /basement/filebox
root@annie:~# screen mv -v /export/* /basement/filebox/
...
root@annie:~# df -k
...
/dev/sdc1        1922728820      512 1922728820    0% /export
...
basement         1787817216 906994176 880823040   51% /basement
...
```

Repartition old drive and add the first partition to the zfs pool as a mirror.

```
root@annie:~# umount /export
... adjust /etc/fstab if necessary ...
root@annie:~# parted /dev/sdc1
(parted) mklabel gpt
Warning: Partition(s) on /dev/sdc are being used.
Ignore/Cancel? I

(parted) mkpart zfs zfs 0% -100512MB
(parted) mkpart efi fat32 -100512MB -100000MB
(parted) mkpart lnx ext2 -100000MB 100%
(parted) set 2 boot on
root@annie:~# reboot

root@annie:~# zpool attach -f basement wwn-0x5000039ff3c899c1-part1 wwn-0x5000039ff3c2ca97-part1
root@annie:~# zpool status
... should show both disks and note (resilvering)
```

Wait for resilvering to finish (1.7T took about 1.75 hours)

```
don@annie:~$ zpool status
pool: basement
state: ONLINE
scan: scrub repaired 0B in 1h44m with 0 errors on Sun Sep  9 02:08:43 2018
```

config:

NAME	STATE	READ	WRITE	CKSUM
basement	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
wnn-0x5000039ff3c899c1-part1	ONLINE	0	0	0
wnn-0x5000039ff3c2ca97-part1	ONLINE	0	0	0

errors: No known data errors

5.109 ZFS IS ALL THE RAGE

Zfs is recommended by the UBUNTU team for LXC/LXD and it has its positives but just like everything else the damned kids fixed is plenty fucking broke. It does not play well with others and it will fuck you in the most subtle and substantial ways.

Look at the disks.

```
root@bs2020:~# fdisk -l
Disk /dev/sda: 136.8 GiB, 146815733760 bytes, 286749480 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x264ef27d

Device      Boot  Start        End  Sectors  Size Id Type
/dev/sda1                2048  286749479  286747432  136.7G 83 Linux

Disk /dev/sdb: 136.8 GiB, 146815733760 bytes, 286749480 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0xc06248cd

Device      Boot  Start        End  Sectors  Size Id Type
/dev/sdb1                2048  85448703   85446656  40.8G 83 Linux
/dev/sdb2          85450750  286748671  201297922   96G  5 Extended
/dev/sdb5          85450752  286748671  201297920   96G 82 Linux swap / Solaris

Disk /dev/sdc: 136.8 GiB, 146815733760 bytes, 286749480 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: 037F919F-B203-449D-A74D-9F285A3B89BD

Device      Start        End  Sectors  Size Type
/dev/sdc1    2048  286749446  286747399  136.7G Linux filesystem

Disk /dev/sde: 136.8 GiB, 146815733760 bytes, 286749480 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: 66247EC5-E595-44A0-B6B8-F4A18179D457

Device      Start        End  Sectors  Size Type
/dev/sde1    2048  286749446  286747399  136.7G Linux filesystem

Disk /dev/sdd: 465.8 GiB, 500107862016 bytes, 976773168 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: C1068A13-6375-48DE-A72B-0F9B3223B6DE

Device      Start        End  Sectors  Size Type
/dev/sdd1    2048  976773134  976771087  465.8G Linux filesystem

root@bs2020:~#
```

There is no connection between this and zfs.

```
root@bs2020:~# zpool status
pool: lxd4dev
state: ONLINE
scan: scrub repaired 0 in 0h4m with 0 errors on Sun Jan 14 00:28:08 2018
config:

    NAME        STATE        READ WRITE CKSUM
    lxd4dev      ONLINE       0     0     0
    sdc1         ONLINE       0     0     0

errors: No known data errors

pool: lxd4infra
state: ONLINE
scan: scrub repaired 0 in 0h0m with 0 errors on Sun Jan 14 00:24:49 2018
config:
```

NAME	STATE	READ	WRITE	CKSUM
lxd4infra	ONLINE	0	0	0
sda1	ONLINE	0	0	0

errors: No known data errors
root@bs2020:~#

if either of these are missing you need to zpool import the disk in its new location. In the long term the disks should be set up to reference their UUIDs but this requires that the pools not be in use (IE Single user mode). This is further complicated by the fact that the pools had to be created in completely different way for lxc and lxd. Hope they fix this in 18.04

5.109.1 Linkdump

Good ones

- <http://kbdone.com/zfs-snapshots-clones/>
- <http://manpages.ubuntu.com/manpages/xenial/man8/zfs.8.html>
- <https://www.howtoforge.com/tutorial/how-to-use-snapshots-clones-and-replication-in-zfs-on-linux/>
-

Fodder

- <https://www.thegeekdiary.com/zfs-tutorials-creating-zfs-snapshot-and-clones/>
- <http://lxd.readthedocs.io/en/latest/backup/#container-backup-and-restore>
- <https://forums.freenas.org/index.php?threads/zfs-send-to-external-backup-drive.17850/>
- <https://www.freebsd.org/cgi/man.cgi?query=zfs>
- <https://www.datto.com/uk/blog/four-ways-to-use-zfs-snapshots>
- <https://forum.proxmox.com/threads/adding-ssd-for-cache-zil-l2arc.25187/>
- <https://www.freebsd.org/cgi/man.cgi?query=zfs>

5.110 kb2018 install bash history.

```
df -k
fdisk -l
apt-get update&&apt-get dist-upgrade&& apt-get autoremove
nano /etc/ssh/authorized_keys
ssh bs2020
ssh bs2020.suspectdevices.com
ssh feurig@bs2020.suspectdevices.com
scp feurig@bs2020.suspectdevices.com:steve.id .
ls .ssh
ls
pwd
exit
apt-get update
apt-get dist-upgrade
apt-get install zfs*
apt-get install bridgeutils
apt-get install bridg*
apt-get install nfs-kernel-server samba-common-bin zfs-initramfs zfs-dracut
ip a
ping archive.ubuntu.com
apt-get install openssh-server
service openssh-server status
service openssh status
service ssh status
ip a
su -feurig
su - feurig
nano /etc/default/grub
update-grub
nano /etc/default/grub
nano /boot/grub/menu.lst
nano /etc/netplan/50-cloud-init.yaml
ip a
nano /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg
nano /etc/netplan/50-cloud-init.yaml
nano /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg
ip a
nano /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg
nano /etc/netplan/50-cloud-init.yaml
netplan apply
nano /etc/netplan/50-cloud-init.yaml
netplan apply
df -k
fdisk -l
ip a
ps -ef
echo FUCKOFF> /dev/ttyS1
reboot
lxd --version
vipw
vigr
useradd -help
useradd -u 1001 -g 1001 -Gwheel,adm,sudo,plugdev,root -m joe -C "Joe Dumoulin"
vigr
useradd -u 1001 -g 1001 -Gwheel,adm,sudo,plugdev,root -m joe -C "Joe Dumoulin"
useradd -u 1001 -g 1001 -Gadm,sudo,plugdev,root -m joe -C "Joe Dumoulin"
useradd -u 1001 -g 1001 -Gadm,sudo,plugdev,root -m -C "Joe Dumoulin" joe
useradd -u 1001 -g 1001 -Gadm,sudo,plugdev,root -m -c "Joe Dumoulin" joe
su - joe
vipw -s
vipw
ls
su - joe
reboot
last
fdisk -l
reboot
fdisk -l
reboot
fdisk -l
fdisk -l|grep Disk
fdisk -l|grep Disk\
fdisk -l|grep Disk\ \
fdisk /dev/sda
apt-get install golang
nano /etc/netplan/50-cloud-init.yaml
netplan apply
nano /etc/netplan/50-cloud-init.yaml
netplan apply
ip a
nano /etc/netplan/50-cloud-init.yaml
netplan apply
ip a
fdisk -l
parted /dev/sdb
fdisk -l
lxd-init
lxd init
```



```

zfs list
ls
ps -ef
ip a
ls
cat bs2020root.id>> /etc/ssh/authorized_keys.save
nano /etc/ssh/authorized_keys.save
cat ~root/.ssh/authorized_keys
cat bs2020root.id>> ~root/.ssh/authorized_keys
tail /var/log/syslog
lxd profile show
lxc profile show
lxc profile show default
ls
lxc profile edit default
cat susdev.yaml
lxc profile edit default
lxc image list
lxc image info
lxc image list ubuntu:*
lxc image alias list ubuntu:*
lxc image alias list ubuntu:* 18.04
lxc image alias list ubuntu:*server* 18.04
lxc image copy ubuntu:18.04 local: --alias ubuntu-lts
lxc launch ubuntu-lts guenter
lxc list
lxc attach guenter
lxc attach exec guenter bash
lxc exec guenter bash
lxc list
ping 192.202.31.134
ip a
lxc exec guenter bash
brctl show
lxc exec guenter bash
nano /etc/sysctl.conf
sysctl -p
lxc exec guenter bash
reboot
apt-get update&&apt-get dist-upgrade
fdisk -l
ls
cat joes.keys
ls
cat ~joe/.ssh/authorized_keys
last|less
ls
zfs list
clear
zfs list
lxc list
networkctl list
nano /etc/netplan/50-cloud-init.yaml
netplan generate
netplan apply
networkctl list
ip a
up br0 up
ip br0 up
ip link br0 up
ip a
ifconfig br0 up
ip a
lxc list
lxc list
lxc stop guenter
lxc list
ip a
ifconfig br1 up
lxc start guenter
ip a
lxc info guenter
lxc exec guenter bash
ps -ef
lxc list
nano /etc/systemd/network/br1.network
nano /etc/systemd/network/br0.network
reboot
ls
lxc list
zpool status
df -k
ip a
lxc shutdown guenter
lxc stop guenter
lxc edit guenter
lxc help
lxc config edit guenter
lxc profile copy default
lxc profile copy default infra
lxc edit profile infra
lxc profile edit infra
lxc delete guenter
networkctl list

```

```

nano /etc/netplan/50-cloud-init.yaml
lxc create ubuntu-lts guenter -p infra
lxc launch ubuntu-lts guenter -p infra
lxc init local:ubuntu-lts guenter -p infra
lxc image list
lxc init local:ubuntu-lts guenter -p infra
lxc list
lxc exec guenter bash
lxc start guenter bash
lxc start guenter
lxc exec guenter bash
lxc list
lxc exec guenter bash
lxc list
lxc image list images:
grep debian
lxc image list images:|grep debian
lxc image list debian:
lxc image remote
lxc image list
lxc image list images:|grep centos
lxc image list images:|grep redhat
lxc image list images:|grep fedora
lxc image list images:|grep suse
lxc list
ps -ef
zfs list
lxc profile edit default
lxc profile show default
lxd init
fdisk -l
zpool list
zpool status
lxd init
zpool status
lxc list
lxd profile edit default
lxc profile edit default
lxc profile edit infra
ls
lxc image list
lxc init ubuntu-lts larry
lxc start larry
lxc list
zpool status
lxc init ubuntu-lts douglas
lxc stop larry
lxc delete larry
lxc start douglass
lxc start douglas
lxc exec douglas bash
lxc list
clear
ip a
lxc config set core.https_address 192.168.31.159:8443
lxc config set core.trust_password w3r3n3t$
lxc remote add kb2018 192.168.31.159
lxc remote list
lxc remote remove kb2018
lxc remote list
lxc profile list
lxc profile copy default susdev
lxc profile list
lxc list
lxc start harvey
lxc exec harvey bash
lxc list
lxc destroy harvey
lxc stop harvey
lxc delete harvey
lxc list
apt-get install htop
htop
ps -ef
df -k
lxc list
zfs list
lxc delete harvey
zfs list
lxc profile delete susdev
lxc profile list
lxc profile copy infra susdev
lxc profile list
lxc info teddy
zfs list
lxc profile delete susdev
lxc start teddy
lxc list
dig digithink.com @dns2.digithink.com
dig digithink.com @dns1.digithink.com
lxc profile rename susdev
lxc profile rename susdev susinfra
lxc info teddy
lxc profile copy default susdev

```

```

lxc list
lxc start sandbox
lxc list
lxc info sandbox
lxc exec sandbox bash
lxc list
ls
lxc list
lxc file put traceback.tgz douglas/home/feurig/
lxc file push traceback.tgz douglas/home/feurig/
lxc exec douglas bash
apt-get update&& apt-get dist-upgrade&& apt-get autoremove
ls
lxc list
lxc start ian
lxc list
lxc start ernest
lxc profile copy default susdev18.04
lxc list
lxc start kurt
lxc list
lxc start morgan
lxc list
lxc delete oldtrac
lxc list
lxc start oldtrac
lxc phillip
lxc start phillip
lxc list
lxc start harvey
lxc delete harvey
reboot
lxc list
lxc info phillip
lxc init local:ubuntu-lts sarina
lxc init local:ubuntu-lts sarina --profile=infra
lxc list profile
lxc profile list
zpool list
lxc stop guenter
ls /var/lib/lxd/storage-pools/infra/containers/guenter/
lxc start guenter
ls /var/lib/lxd/storage-pools/infra/containers/guenter/
lxc help
lxc storage help
lxc storage show
lxc storage show infra
lxc help
ls
lxc list
lxc start naomi
lxc put naomiroot.tgz naomi:/home/feurig/
lxc file put naomiroot.tgz naomi:/home/feurig/
lxc file push naomiroot.tgz naomi:/home/feurig/
lxc file push naomiroot.tgz naomi:/home/feurig/
lxc exec naomi bash
tar -xzf naomiroot.tgz
pwd
ls
cd /var/lib/lxc/naomi/
ls -ls
ls rootfs/etc/
ls -ls rootfs/etc/
df -k
ls /var/lib/lxd/storage-pools/infra/containers/naomi/
mv /var/lib/lxd/storage-pools/infra/containers/naomi/rootfs /var/lib/lxd/storage-pools/infra/containers/naomi/
mv rootfs /var/lib/lxd/storage-pools/infra/containers/naomi/
lxc list
lxc help
lxc exec naomi bash
ls
lxc list
lxc exec naomi bash
lxc config set naomi security.privileged true
lxc exec naomi bash
lxc list
fdisk -l
cd /srv/installmedia/
ls
wget https://downloads.sourceforge.net/gparted/gparted-live-0.32.0-1-amd64.iso
wget
wget https://download.fedoraproject.org/pub/fedora/linux/releases/28/Server/x86_64/iso/Fedora-Server-28-1-20180814-1-1-amd64.iso
ls
wget https://cdimage.debian.org/debian-cd/current/amd64/iso-cd/debian-9.5.0-amd64-netinst.iso
ls
wget https://saimei.ftp.acc.umu.se/debian-cd/current/amd64/iso-dvd/debian-9.5.0-amd64-DVD-2.iso
ip a
apt-get install samba samba-common-bin
apt-get install samba samba-common-bin
mkdir /srv/installmedia
cd /srv/installmedia/
wget http://releases.ubuntu.com/18.04.1/ubuntu-18.04.1-live-server-amd64.iso?_ga=2.217616086.1525111111.1525111111.1525111111
apt-get remove samba
apt-get autoremove

```

```

ls
mv ubuntu-18.04.1-live-server-amd64.iso\?_ga\=2.217616086.1525765299.1538535940-781929701.1526740
ls
apt-get install nfs-kernel-server
ls -ls
nano /etc/exports
exportfs -a
nano /etc/exports
exportfs -a
showmount -e
useradd -c"nfs client" nfs
passwd nfs
ls
ip a
tail /var/log/syslog
tail -f /var/log/syslog
nano /etc/exports
exportfs -a
showmount -e
ls -ls /srv/installmedia/ubuntu-18.04.1-live-server-amd64.iso
showmount -e
passwd nfs
su - feurig
vipw
su - feurig
ls
showmount -e
ip -a
ip a
showmount -e
ufw
ufw help
ufw status
showmount -e
nano /etc/exports
ls
chown -R nfs /srv/installmedia
ls -ls
wget https://cdimage.debian.org/debian-cd/current/amd64/iso-dvd/debian-9.5.0-amd64-DVD-1.iso
ls
arp -a
ssh feurig@192.168.31.200
ssh feurig@192.168.31.158
ls
cd /srv/installmedia/
ls
cd /srv/installmedia/
ls
ls -ls
userdel nfs
ls
ls-ls
ls -ls
chown root *
ls -ls
ls
arp -a
ping 192.168.31.196
ping 192.168.31.200
ssh feurig@192.168.31.200
ping 192.168.31.158
ssh feurig@192.168.31.158

```

5.111 Unifi install

```
# ----- roles/ubuntu-unifi-server/tasks/main.yml
- name: download unifi-latest script
  get_url:
    url: https://get.glennr.nl/unifi/install/install_latest/unifi-latest.sh
    dest: /root/unifi-latest.sh
    mode: '0700'

- name: Run Easy Unifi Script.
  command: "bash /root/unifi-latest.sh --skip --add-repository"

- name: install nginx
  apt:
    state: latest
    name: nginx

- name: remove default nginx site
  file:
    path: /etc/nginx/sites-enabled/default
    state: absent

- name: Seed nginx configuration
  copy:
    src: "../files/nginx.conf"
    dest: /etc/nginx/nginx.conf
#
- name: Restart nginx service
  service:
    name: nginx
    enabled: yes
    state: restarted
```

```
# ----- ansible/roles/ubuntu-unifi-server/files/nginx.conf
# ---- simplest redirection I could figure out

user www-data;
worker_processes auto;
pid /run/nginx.pid;

include /etc/nginx/modules-enabled/*.conf;

events {}

stream {
    upstream unifi {
        server localhost:8443;
    }
    server {
        listen      443;
        proxy_pass  unifi;
    }
}

http {
    server {
        listen 80 default_server;
        listen [::]:80 default_server;
        server_name _;
        return 301 https://$host:8443$request_uri;
    }
}
```