

Creating OpenIndiana EC2 image

Install OpenIndiana as paravirtualized domU using text installer.

I have used Debian 6 64-bit with Xen 4.1 as a host. Sample config file:

```
memory = 1024
name = 'oi'
vcpus = 4
vif = [ 'mac=00:16:3e:48:a6:54' ]
disk = [ 'file:/root/oi-dev-147-text-x86.iso,xvda:cdrom,r', 'file:/vms/oi/main.raw,xvdb,w' ]
bootloader = '/usr/bin/pygrub'
```

Fix the /rpool/boot/grub/menu.lst.

Change /platform/i86pc/kernel/\$ISADIR/unix to /platform/i86xpv/kernel/\$ISADIR/unix in the menu.lst of the installed system.

Create /dev/{disk,rdisk} symbolic links.

When booted in EC2 environment, the first disk is represented by /devices/xpvd/xd@2048, the second one by /devices/xpvd/xd@2064 and so on. Links in /dev/dsk and /dev/rdisk pointing at /devices/xpvd/xd* are required for the system to boot. They need to be present on the active zfs root dataset. Example can be found at the bottom of this page [1].

To simulate disk setup in EC2 environment 'disk' line as below can be put in the Xen cfg file.

```
disk = [ 'file:/vms/oi/pvgrub64.img,2048,w', 'file:/vms/oi/main.raw,2064,w' ]
```

Note that devfsadm is not happy about the xdf@ numbers and prints out some harmless error messages at boot (the numbers are probably out of range). This is probably the reason why the text installer is unable to detect disks even after pre-creating the links using the 'Shell' feature of the install cd as well.

Prepare zfs enabled pv-grub boot disk image.

Create fat16 or fat32 filesystem on unpartitioned disk image. Disk size smaller or equal 1GB is recommended. Minimal EBS disk allocation is 1GB by the way. 16MB fat16 disk image worked ok for me.

Create /boot/grub/menu.lst file for Amazon's pv-grub.

```
# cat /mnt/boot/grub/menu.lst
timeout 0
default 0
title chainload pv-grub
root (hd0)
kernel /pv-grub.gz (hd0)/menu.lst
```

Create /menu.lst file for custom zfs enabled pv-grub.

```
# cat /mnt/menu.lst
timeout 0
default 0
title load /rpool/boot/grub/menu.lst
findroot (pool_rpool,0,a)
configfile /boot/grub/menu.lst
```

Copy 32 or 64-bit pv-grub binary to /pv-grub.gz.

pv-grub is part of official Xen source tree. It can be built by going to xen-unstable.hg/stubdom/ directory and running 'make' or 'XEN_TARGET_ARCH=x86_32 make' for 32-bit version. Resulting files can be found in xen-unstable.hg/stubdom/mini-os-x86_{32,64}-grub/mini-os.gz. For zfs support attached grub.patch patch needs to be applied first. Also 999sun-grub-addons.patch needs to be copied to xen-unstable.hg/stubdom/grub.patches. Files are attached.

Upload disk images to EC2 EBS storage (disk image pool).

I couldn't find a way to upload disk image from local machine directly to EBS. I had to attach newly created blank EBS volumes to running EC2 instance (for example temporary Ubuntu one) and then dd my prepared images over. You can attach volumes at run-time which is quite handy.

After running dd I took snapshots of the images. Elasticfox Firefox extension is great for working with disk images/snapshots or EC2 in general by the way.

Create OpenIndiana AMI (VM template).

After installing EC2 cli tools on my local machine I ran:

```
ec2reg -a x86_64 -b /dev/sda=snap-id-of-the-pvgrub-image -b /dev/sda=snap-id-of-the-main-oi-image --kernel aki-427d952b -n OpenIndiana-64
```

aki-427d952b is the kernel image id of the pv-grub 64-bit provided by Amazon.

Start the instance based on the newly created AMI.

Initially the instance will not have any network connectivity. I have not investigated why exactly. Missing device links perhaps. Waiting for the instance to start and then rebooting it is a workaround. To check the status of the VM it is possible to look at the console output either by using Elasticfox or cli tools. The output does not appear immediately. There is typically few minute delay.

[1] Device links:

```
oi@openindiana:~$ ls -l /dev/dsk /dev/rdisk
/dev/dsk:
total 21
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0p0 -> ../../devices/xpvd/xdf@2048:q
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0p1 -> ../../devices/xpvd/xdf@2048:r
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0p2 -> ../../devices/xpvd/xdf@2048:s
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0p3 -> ../../devices/xpvd/xdf@2048:t
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0p4 -> ../../devices/xpvd/xdf@2048:u
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s0 -> ../../devices/xpvd/xdf@2048:a
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s1 -> ../../devices/xpvd/xdf@2048:b
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s10 -> ../../devices/xpvd/xdf@2048:k
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s11 -> ../../devices/xpvd/xdf@2048:l
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s12 -> ../../devices/xpvd/xdf@2048:m
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s13 -> ../../devices/xpvd/xdf@2048:n
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s14 -> ../../devices/xpvd/xdf@2048:o
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s15 -> ../../devices/xpvd/xdf@2048
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s2 -> ../../devices/xpvd/xdf@2048:c
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s3 -> ../../devices/xpvd/xdf@2048:d
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s4 -> ../../devices/xpvd/xdf@2048:e
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s5 -> ../../devices/xpvd/xdf@2048:f
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s6 -> ../../devices/xpvd/xdf@2048:g
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s7 -> ../../devices/xpvd/xdf@2048:h
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s8 -> ../../devices/xpvd/xdf@2048:i
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t0d0s9 -> ../../devices/xpvd/xdf@2048:j
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0p0 -> ../../devices/xpvd/xdf@2064:q
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0p1 -> ../../devices/xpvd/xdf@2064:r
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0p2 -> ../../devices/xpvd/xdf@2064:s
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0p3 -> ../../devices/xpvd/xdf@2064:t
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0p4 -> ../../devices/xpvd/xdf@2064:u
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s0 -> ../../devices/xpvd/xdf@2064:a
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s1 -> ../../devices/xpvd/xdf@2064:b
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s10 -> ../../devices/xpvd/xdf@2064:k
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s11 -> ../../devices/xpvd/xdf@2064:l
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s12 -> ../../devices/xpvd/xdf@2064:m
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s13 -> ../../devices/xpvd/xdf@2064:n
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s14 -> ../../devices/xpvd/xdf@2064:o
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s15 -> ../../devices/xpvd/xdf@2064
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s2 -> ../../devices/xpvd/xdf@2064:c
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s3 -> ../../devices/xpvd/xdf@2064:d
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s4 -> ../../devices/xpvd/xdf@2064:e
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s5 -> ../../devices/xpvd/xdf@2064:f
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s6 -> ../../devices/xpvd/xdf@2064:g
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s7 -> ../../devices/xpvd/xdf@2064:h
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s8 -> ../../devices/xpvd/xdf@2064:i
lrwxrwxrwx 1 root root 29 2011-04-06 01:03 c0t1d0s9 -> ../../devices/xpvd/xdf@2064:j
/dev/rdisk:
total 21
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0p0 -> ../../devices/xpvd/xdf@2048:q,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0p1 -> ../../devices/xpvd/xdf@2048:r,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0p2 -> ../../devices/xpvd/xdf@2048:s,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0p3 -> ../../devices/xpvd/xdf@2048:t,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0p4 -> ../../devices/xpvd/xdf@2048:u,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s0 -> ../../devices/xpvd/xdf@2048:a,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s1 -> ../../devices/xpvd/xdf@2048:b,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s10 -> ../../devices/xpvd/xdf@2048:k,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s11 -> ../../devices/xpvd/xdf@2048:l,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s12 -> ../../devices/xpvd/xdf@2048:m,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s13 -> ../../devices/xpvd/xdf@2048:n,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s14 -> ../../devices/xpvd/xdf@2048:o,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s15 -> ../../devices/xpvd/xdf@2048 ,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s2 -> ../../devices/xpvd/xdf@2048:c,raw
```

```
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s3 -> ../../devices/xpvd/xdf@2048:d,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s4 -> ../../devices/xpvd/xdf@2048:e,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s5 -> ../../devices/xpvd/xdf@2048:f,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s6 -> ../../devices/xpvd/xdf@2048:g,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s7 -> ../../devices/xpvd/xdf@2048:h,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s8 -> ../../devices/xpvd/xdf@2048:i,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t0d0s9 -> ../../devices/xpvd/xdf@2048:j,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0p0 -> ../../devices/xpvd/xdf@2064:q,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0p1 -> ../../devices/xpvd/xdf@2064:r,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0p2 -> ../../devices/xpvd/xdf@2064:s,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0p3 -> ../../devices/xpvd/xdf@2064:t,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0p4 -> ../../devices/xpvd/xdf@2064:u,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s0 -> ../../devices/xpvd/xdf@2064:a,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s1 -> ../../devices/xpvd/xdf@2064:b,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s10 -> ../../devices/xpvd/xdf@2064:k,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s11 -> ../../devices/xpvd/xdf@2064:l,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s12 -> ../../devices/xpvd/xdf@2064:m,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s13 -> ../../devices/xpvd/xdf@2064:n,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s14 -> ../../devices/xpvd/xdf@2064:o,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s15 -> ../../devices/xpvd/xdf@2064 ,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s2 -> ../../devices/xpvd/xdf@2064:c,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s3 -> ../../devices/xpvd/xdf@2064:d,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s4 -> ../../devices/xpvd/xdf@2064:e,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s5 -> ../../devices/xpvd/xdf@2064:f,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s6 -> ../../devices/xpvd/xdf@2064:g,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s7 -> ../../devices/xpvd/xdf@2064:h,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s8 -> ../../devices/xpvd/xdf@2064:i,raw
lrwxrwxrwx 1 root root 33 2011-04-06 01:03 c0t1d0s9 -> ../../devices/xpvd/xdf@2064:j,raw
```