

Mirroring OpenIndiana

Data available for mirroring

There are two separate groups of data available for you to mirror:

- Downloadable Content (DLC) means the ISO and USB installation images, plus source archives used by developers to build the userland software available in OpenIndiana. The content is accessible through dlc.openindiana.org.
Size: **25 GB** (as of 2012-07-11).
- Packages Files (pkgfiles) is a much simplified version of the package repository running at pkg.openindiana.org, which only contains the actual data from the packages, and no meta data at all. This allows you to host a mirror for the `pkg(1)` client with a static web server such as Apache, lighttpd, nginx, etc., without having to worry about `pkg.depotd(1M)`. It also means you don't have to run the mirror on OpenIndiana, if you already have a mirror server set up with another operating system.
Size: **17 GB** (as of 2012-07-11).

Mirror infrastructure

Requirements

We try to keep the requirements low for becoming a mirror in order to make it easier for people to contribute to the project. The requirements are:

1. The mirror must be run on a machine running 24/7.
2. It should at least have a 100 Mbit Internet uplink.
3. The mirror must check for updates to the files once or twice a day.
4. The web server software must serve a file named "index.html" if a directory path is requested, and the file exists.
5. The web server must show a directory listing if a directory path is requested, and no "index.html" file is found.
6. Symbolic links on the file system must be followed by the web server.

The last three points can for instance be fulfilled by an Apache HTTPD with the following settings configured:

```
DirectoryIndex index.html
<Directory />
    Options -All +Indexes +FollowSymlinks
    IndexOptions +FancyIndexing +NameWidth=*
</Directory>
```

And the following is comparable for Nginx:

```
server {
    location / {
        index index.html;
        autoindex on;
    }
}
```

Getting the data

After your mirror is configured

Old content follows

This page covers how to set up an OpenIndiana b147 machine for mirroring the data currently available on pkg.oi.o/dev and dlc.oi.o, it should however work fine on other versions of OpenSolaris (post b134) and OpenIndiana. It should also work without problems within a zone.

The guide will not cover how to install OpenIndiana, have a look at the [2. Installing or Upgrading](#) guide instead. You should also install `rsync` before starting on the guide, it can be installed by executing `pkg install rsync`.

`wget` will be used to download files within the guide. It is not installed per default in a new zone but can be installed by installing the `wget` package, do however feel free to use another download utility such as `curl` if you prefer to do so.

All commands in the guide are expected to be run with root permissions, either by using `pexec`, `sudo` or simply a shell with root permissions.

Around 2GB is needed for `/dev`, around 7GB for `/legacy`, and currently around 6GB for `dlc.o.i.o`.

Getting the files

There are two parts to getting the files needed for mirroring; the repository data and download content. We will start off with the repository data.

Repository data

You should start off with the newest available tarball available on the [DLC](#), next we are going to use `rsync` for getting the repository data updated. Here I am the putting the data for the repository into `/export/pkg/dev/` with ZFS compression disabled (the depot files are already zipped):

```
zfs create -p -o atime=off -o compression=off rpool/export/pkg/dev
cd /export/pkg/dev/
wget http://dlc-int.openindiana.org/repos/oi_151a.tar.bz2
gtar -jxf oi_151a.tar.bz2
mv dev_151a/* ./
rmdir dev_151a/
rsync -a pkg-origin.openindiana.org::pkgdepot-dev /export/pkg/dev/
```

That leaves us with a complete updated `pkg` root inside `/export/pkg/dev/`, so we move onto the DLC data.

DLC content

The DLC data is pretty static, so we are going to fetch it using `rsync`, and place the data inside `/export/dlc/`:

```
zfs create -p -o atime=off -o compression=off rpool/export/dlc
rsync -a dlc-origin.openindiana.org::dlc /export/dlc/
```

We'd recommend only mirroring the `isos` directory if you're short on disk space.

And that's it, those are the files needed to serve clients. Next we are going to set up a couple of services for letting other people get the files.

Serving the data

DNS

We have a DNS convention as follows:

```
Download Area:  dlc-1.gb.openindiana.org   = Download server 1 for GB (Great Britain)
Package Server: pkg-1.gb.openindiana.org   = Package server 1 for GB (Great Britain)
```

So for example if you're in France, instead of `gb` it would be `fr`.

Please contact Alasdair Lumsden ([alasdairrr at gmail](mailto:alasdairrr@gmail.com)) to request an allocated DNS entry. Please provide information regarding your mirror (Your name, contact details, who is sponsoring the mirror, where the mirror is geographically located and the bandwidth available to the mirror) and your DNS entry will be created and the mirror details added to the mirror list.

pkg.depotd



Warning

You need to serve the files with `pkg.depotd` from `oi_147` or later. If you're not running `oi_147` you can compile and install the latest `pkg5`.

We are going to start with setting up the PKG repository, since that's quite easy. We only need to configure the a SMF service instance of `svc:/application/pkg/server` to match the directory to serve files from and the port number to run as:

```
svccfg -s pkg/server
add dev
select dev
addpg pkg application
setprop pkg/inst_root = astring: "/export/pkg/dev/"
setprop pkg/port = count: 10002
exit
svcadm refresh pkg/server:dev
svcadm enable pkg/server:dev
```

That should leave you with a pkg repository running on port 10002 on the machine, try to browse it with a web browser to make sure it works before we move on.

Theming the website

In order to make the web interface for the repository look the part, you should use the OpenIndiana theme available. This is easily installed through `rsync`:

```
rsync -a pkg-origin.openindiana.org::pkgdepot-branding/web/ /usr/share/lib/pkg/web/
```

After that you simply have to restart the web server with `pkg restart pkg/server:dev` and the new theme is online.

Apache

We are going to serve all the data through Apache, it will make it easy to serve multiple repositories on the same IP address, but with different hostnames or paths. Start by installing Apache on the system:

```
pkg install apache-22
```

Here is an example Apache VirtualHost for serving the pkg repos taken from the pkg-origin server:

```

<VirtualHost *:80>
    ServerName pkg.openindiana.org
    ServerAlias pkg-origin.openindiana.org

    DocumentRoot /var/www/htdocs
    CustomLog /var/www/logs/access_log common
    ErrorLog /var/www/logs/error_log

    <Directory /var/www/htdocs>
        Order allow,deny
        Allow from all
    </Directory>

    <Directory /export/pkg>
        Order allow,deny
        Allow from all
        Options -All
    </Directory>

    SetOutputFilter DEFLATE

    MaxKeepAliveRequests 10000
    ProxyTimeout 60
    ProxyRequests Off
    AllowEncodedSlashes On

    SetEnvIf Request_URI .*\/file\/[01]\/.* no-gzip no-vary

    <LocationMatch ".*\/catalog.attrs">
        Header set Cache-Control no-cache
    </LocationMatch>

    RewriteEngine On

    RewriteRule ^\/dev$ /dev/ [R]
    RewriteRule ^\/dev\/file\/[01]\/(..)(.*) /export/pkg/dev/file\/$1\/$1$2
    ProxyPass /dev/ http://127.0.0.1:10002/ nocanon disablereuse=On

    RewriteRule ^\/legacy$ /legacy/ [R]
    RewriteRule ^\/legacy\/file\/[01]\/(..)(.*) /export/pkg/legacy/file\/$1\/$1$2
    ProxyPass /legacy/ http://127.0.0.1:10134/ nocanon disablereuse=On
</VirtualHost>

```

For serving dlc, a standard configuration should be sufficient.

Write how to modify pkg/proxy_base