



Porting HUBzero to RHEL6

Edward Duffy, Ph.D.
Computational Scientist, CITI
Clemson University

Sep 24, 2012

Outline

- Why undergo a port?
- Technical issues porting Debian packages to Red Hat packages
- Questions

Purpose

- Port the HUBzero platform to Red Hat Enterprise Linux (version 6)
- Currently, HUBzero is distributed in an `apt-get` repository¹
 - All configuration scripts assume a Debian host
- We are attempting to repackaging everything as `.rpm` packages
 - Stand-up a `yum` repository to host packages
 - Update configuration scripts

¹<http://packages.hubzero.org/deb/pool/main/h>

Why Bother?

- Clemson manages a Red Hat-centric datacenter
- University pays for professional support from Red Hat, Inc.
- Staff is trained on Red Hat supported software, and responsible for 24/7 maintenance of the servers in the data center

Why bother? (cont'd)

- For reliability, the servers they support should fit in with their workflow:
 - OS installation
 - Network configuration
 - Security auditing and upgrades
 - Firewall rules
 - Backup schedules & data recovery
 - Virtual Machine management
 - System health and continued monitoring

System software differences

- Whenever possible, we want to stick with the version from the official repositories
 - Easier on the support staff
- Only repackage non-H0 software when incompatibilities arise
 - Although it would be preferable to patch the H0 source

Python

- RHEL6 supplies Python 2.6
 - H0 has a dependence on Python 2.5
 - Python makes it easy to install multiple versions side-by-side (make `altinstall`)
 - Repackaged non-standard Python libraries, incl: ldap, mysql, openssl)

OpenLDAP

- RHEL6 requires OpenLDAP 2.4.23
 - The yum package manager will break with a lower version
 - The configuration is vastly different from Debian's OpenLDAP 2.4.12
 - Debian configuration file: `/etc/openldap/slapd.conf`
 - RHEL configuration directory: `/etc/openldap/slapd.d`

OpenLDAP (cont'd)

- Configuration defined by .ldif files:

```
% find /etc/openldap/slapd.d/  
/etc/openldap/slapd.d/  
/etc/openldap/slapd.d/cn=config.ldif  
/etc/openldap/slapd.d/cn=config  
/etc/openldap/slapd.d/cn=config/cn=schema.ldif  
/etc/openldap/slapd.d/cn=config/olcDatabase={0}config.ldif  
/etc/openldap/slapd.d/cn=config/cn=schema
```

and so on...

Smaller issues

- Apache is the same version
 - Although the distributions use different user name and group names
- For the maxwell service, keeping the Debian-based containers
- `start-stop-daemon` taken from `dpkg` to get startup scripts working

RPMs

- Currently, around 40 .rpms
- Some are partial, only the data portion has been ported
 - Post-installation configuration still needs work (more on this in a few slides)
- Close to the original sources as possible
 - Some changes for configuration or compiling needed
 - Any patching is performed inside the .spec file
- Even when RPMs complete, the big hz-install script still needs to be ported to RHEL

Issues with porting packages (by example)

- HUBzero packages are not available in the traditional tarball format
- Packages are stored as .deb
- There's no familiar `./configure && make && make install`
- .debs must be unpacked, and repackaged as .rpms

Repackaging

- Take hubzero-config-1.0.1, for example
- Two main sub-packages: control.tar.gz and data.tar.gz

```
% ls
```

```
hubzero-config_1.0.0-1_all.deb
```

```
% ar x hubzero-config_1.0.0-1_all.deb
```

```
% ls -l
```

```
control.tar.gz
```

```
data.tar.gz
```

```
debian-binary
```

```
hubzero-config_1.0.0-1_all.deb
```


Repackaging

- control.tar.gz contains the package metadata and scripts fired at predetermined times

```
% tar zxvf control.tar.gz
./
./templates    [M]
./control      [M]
./preinst      [S]
./md5sums      [M]
./postrm       [S]
./config       [S]
./postinst     [S]
```


Repackaging

- dkpg allows the package maintainer to define install-time variables
 - Variables may be given sensible defaults, the user may be prompted to provide values, and are query-able from all scripts
- For example, the variable storing the fully qualified domain name, fqdn

template

- Variables are defined in the template file, with name, type, description, and prompt

Template: hubzero-config/fqdn

Type: string

Description: FQDN for hub:

Enter the FQDN for the system using the format of
fully.qualified.domainname

config

- Here, fqdn is given a default value, added to dpkg's database

```
db_title "Configuring hubzero-config"
domain=`hostname -f`
db_get hubzero-config/fqdn
if [ -z "$RET" ]; then
    db_set hubzero-config/fqdn $domain
fi
```

- The value may be overridden by the user

```
case "$state" in
    fqdn)
        db_input critical hubzero-config/fqdn || true
```


postinst

- Finally, it's value used to generate the hub's configuration file

```
db_get hubzero-config/fqdn  
fqdn="$RET"
```

```
# ...
```

```
echo "[default]" >> $conffile  
echo "site=$sitename" >> $conffile  
echo "" >> $conffile  
echo "[$sitename]" >> $conffile  
echo "DocumentRoot=/www/$sitename" >> $conffile  
echo "HubName=$sitename" >> $conffile  
echo "HubHost=$fqdn" >> $conffile
```


Package data

- The `data.tar.gz` contain the files that will be installed on the system

```
% tar zxvf data.tar.gz
./
./usr/
./usr/share/
./usr/share/hubzero-config/
./usr/share/hubzero-config/hz-install.tmpl
./usr/share/doc/
./usr/share/doc/hubzero-config/
./usr/share/doc/hubzero-config/copyright
./usr/share/doc/hubzero-config/changelog.Debian.gz
./usr/share/doc/hubzero-config/README.Debian
```


Questions

- Any interest from other RHEL-based universities that host their own hubs?
- Any interest from the HUBzero team to support RHEL installs in parallel with Debian?
 - Afraid of this “port” being out-of-sync with HUBzero-proper, and becoming a “fork”