

Perl in RPM-Land

Dave Cross
Magnum Solutions Ltd
<http://mag-sol.com/>
dave@mag-sol.com

MiltonKeynes.pm
8th September 2008

What is RPM?

- RPM is the “RPM Package Manager”
- Previously “RedHat Package Manager”
- Package manager for Red Hat Linux
- Now used by many distributions
 - RHEL and Centos
 - Fedora
 - SUSE
 - Mandriva
- This talk will concentrate on Fedora and Centos

Distributing RPMs

- An RPM is distributed as a single file
 - `package_name-version.rpm`
- Also a source RPM
 - `package_name-version.src.rpm`
 - Contains source and build instructions
- Install with command line program
 - `rpm -i package_name-version.rpm`
- Or, more usually,
 - `rpm -Uvh package_name-version.rpm`

Problems With rpm

- Dependency hell
- Each RPM contains a list of required RPMs
- Won't install if requirements are missing
- Warns of missing requirements
- But won't install them
- Better to use yum instead

yum

- YellowDog Updater Modified
- `yum install package-name`
- No need to specify version
- Installs latest version
- Traces dependency tree
- Downloads and installs all requirements
- A bit like `apt-get`

Advantages of RPMs

- Dependency resolution
- Easy distribution
- Standard installation method
- Uninstall simply (`rpm -e`)
- Built-in versioning
- Command line tools
- Tools to query RPM database

What About the CPAN Shell?

- The CPAN shell can do most of this
- But it has some problems
- Sysadmins often prefer RPMs
- RPM can handle non-CPAN dependencies
- Other RPMs might depend on CPAN RPMs
- CPAN doesn't update the RPM database
- The siteperl / vendorperl incompatibility

The siteperl / vendorperl Incompatibility

- Standard CPAN installations put modules in the siteperl directory
- RPM standard is to put modules in the vendorperl directory
- vendorperl is before siteperl in @INC
- Therefore mixing RPM-installed modules and CPAN-installed modules is a bad idea
- e.g. Updating an RPM installed module using CPAN

Dave's Rule of Mixing RPMs with CPAN Modules

- **Never mix CPAN-installed modules and RPM-installed modules in the same Perl installation**
- I've been following this rule for about three years now and it has made my life much happier

Corollary

- If you want to have a Perl installation which uses CPAN-installed modules then have one which is completely separate
- The system Perl already comes with a lot of RPM-installed modules so you shouldn't use CPAN to install modules there

RPMs of CPAN Modules

- There are two ways to get hold of RPMs of CPAN modules
- Find a decent RPM repository
- Build your own
- We'll look at both of these

Finding RPMs

- RPMs are usually distributed from online repositories
- yum is configured to work with these repositories
- List available Perl RPMs
 - `yum list available 'perl-*`
- Install chosen RPM
 - `yum install perl-Calendar-Simple`

Yum Repositories for Fedora

- Many Perl RPMs in the default Fedora repository
- Over 900 packages available
- More packages added all the time
- Fedora Perl developers mailing list
 - <https://www.redhat.com/mailman/listinfo/fedora-perl-devel-list>

Other Fedora Repositories

- Livna - <http://rpm.livna.org/>
 - Concentrates on packages that can't be officially distributed by Fedora
 - Only one Perl package (perl-MythTV)
- ATrpms - <http://www.atrpms.net/>
- Dag Wieers -
<http://dag.wieers.com/rpm/>
- Dries Verachtert -
<http://dries.ulyssis.org/rpm/>

Yum Repositories for Centos

- As with Fedora, many modules part of the default Centos repository
- Also uses Fedora packages
 - But generally a few months behind

Other Centos Repositories

- EPEL - <http://fedoraproject.org/wiki/EPEL>
 - Extra Packages for Enterprise Linux
 - Repackaging of Fedora packages
 - Hosted by Fedora project
- RPMForge - <http://rpmforge.net/>
 - Combined repository
 - Dag, Dries and a couple of smaller repositories

Latest Repository News

- Dag, Dries and a couple of smaller repositories have joined forces to form RPMForge
 - <http://rpmforge.net/>
 - Currently supports Centos but not Fedora
- RPM Repo - <https://rpmrepo.org/>
 - Will eventually will combine RPMs from ATrpms, Dag, Dries and others

Building RPMs from CPAN

- I put off doing this for a long time
- Thought it would be complicated
- Turned out to be easier than I imagined
- Good tools help

Building RPMs

- You'll need an RPM build environment
- `yum install rpm-build`
- Create `~/.rpmmacros`
- `%packager Dave Cross <dave@mag-sol.com>`
`%vendor Magnum Solutions`
`%_topdir /home/dave/rpm`
- Create working directories
- `mkdir -p ~/rpm/BUILD ~/rpm/RPMS \`
`~/rpm/SOURCES ~/rpm/SPECS ~/rpm/SRPMS`

Building an RPM

- An RPM file is built from a source distribution and a spec file
- Spec file controls how the RPM is built
- Pre- and post-build processing
- Calendar-Simple-1.20.tar.gz
+ perl-Calendar-Simple.spec
= perl-Calendar-Simple-1.20.1.rpm

Spec Files

- Spec files are a bit of a dark art
 - To me, at least
- Many options controlling the build
- In the general case, constructing a spec file is complex
- But CPAN modules use a standard build mechanism
 - Usually
- So building a spec file can be automated

cpanspec

- cpanspec is the standard tool for creating RPM spec files from CPAN distributions
 - In the Fedora/CentOS world, at least
- `yum install cpanspec`
- `cpanspec module_name`
- `cpanspec Calendar::Simple`
- Many options
- Rarely used
 - In my experience

Building Your RPM

- `cpanspec` will build the RPM for you
- `cpanspec -b cpan-package-name`
- But that gets messy if the spec file needs tweaking
- I use `rpmbuild` to build my RPMs
- `rpmbuild -ba specfile`
- `rpmbuild -ba perl-Calendar-Simple.spec`
- Creates RPM and source RPM
- Then just install using `rpm`

From CPAN to RPM

- `cd ~/rpm`
- `cpanspec Calendar::Simple`
- `mv Calendar-Simple-1.20.tar.gz SOURCES`
- `mv perl-Calendar-Simple.spec SPECS`
- `rpmbuild -ba SPECS/perl-Calendar-Simple.spec`
- `rpm -Uvh /home/dave/rpm/RPMS/noarch/perl-Calendar-Simple-1.20-1.fc9.noarch.rpm`

Common Errors

- All tests must pass
 - Submit any patches back to CPAN author
- Missing files
 - cpanspec often misses installed programs
 - Fix %files section in spec file
- Missing dependencies
 - Generated from PREREQS in Makefile.PL
 - Fix and report to author

Sharing Your Work

- If you build an RPM for a module please consider sharing it
- Put it on a web page
 - <http://rpm.mag-sol.com/>
- Create a real yum repository
- Submit it to the Fedora project
 - Licensing issues
 - <http://fedoraproject.org/wiki/Packaging/Perl>
 - <http://fedoraproject.org/wiki/SIGs/Perl>

Something More Ambitious

- Many CPAN spec files can be generated automatically with `cpanspec`
- Many others can be generated, hand-tweaked and reused for subsequent releases
- The Fedora project currently packages ~10% of CPAN
- Can we help them improve that?

A Rough Plan

- <http://rpm.cpan.org/>
- Monitor CPAN uploads
- Generate automatic specs for new modules
- If the build works distribute the RPM
- If not, sent it to a volunteer for investigation and fixing
- Have as many up to date CPAN RPMs available as possible
- For as many RPM-based Linuxes as possible

A Project

- I mentioned that idea at YAPC::Europe
- People didn't laugh
- People got quite interested
 - Especially Jos
 - <http://debs.pkgs.cpan.org/>
 - His employers are moving to Centos
- A lot of interest in taking it further

Project Status

- A mailing list
 - <http://lists.dave.org.uk/mailman/listinfo/cpanrpm>
 - (Google for 'cpanrpm')
- CPANPLUS::Dist::Fedora
 - Uses CPANPLUS packaging mechanism
- <http://rpm.mag-sol.com> is a Yum repository
- <http://rpm.pkgs.cpan.org/>
 - Currently empty

Next Steps

- Get automatic builds running on <http://rpm.mag-sol.com/>
- Monitor errors
 - And fix
- Add non-CPAN dependencies to specs
- Gradually add more and more CPAN distributions
- Move to <http://rpm.pkgs.cpan.org/>
- Profit!

Questions

- Any questions?
- Any volunteers?