

RDO and CentOS

Ideas for more collaboration

http://www.pixelbeat.org/talks/rdo_centos/

Pádraig Brady – Red Hat

Agenda

- OpenStack Overview
- RDO Overview
- Thoughts for more CentOS collaboration

About Me

Open Source user for 13 years

Python user for 11 years

Gnu coreutils maintainer for 5 years

Red Hatter for 2 years

OpenStack Nova core member for 1.5 years

RDO packaging lead for 1.5 years

What is OpenStack

- IaaS platform (like AWS)
- Open Source
 - unlike VMware, AWS, Compute Engine etc.
 - No closed components
- Leverages lots of existing Linux technologies
- Written mainly in `2.6 <= python < 3.0`
- 3 years old
- Very large project

Havana Release Stats

(6 month release cycle)

- [bitergia havana analysis](#)
- Commits: 13,624 (out of 47K total)
 - Excluding merge commits and bots etc.
 - About 3 per hour
- Source code reviews: 21,228
- Developers: 923
- Organizations: 150

OpenStack is still accelerating

- 10K proposed commits in the last 60 days
- New OpenStack projects in the pipeline
 - Savanna (hadoop)
 - Marconi (SQS, SNS like)
 - Trove (DBaaS)
 - Ironic (baremetal provisioning)
 - Tripleo (OpenStack on OpenStack)

Python projects used (Icehouse milestone 2)

```
$ curl https://raw.githubusercontent.com/openstack/requirements/master/global-requirements.txt |  
sed '/^ *#/d; /^ *$/d; s/ \?[>=<#].*//; /python-.*client/d' | pr -T6 -w80 | expand
```

alembic	greenlet	os-collect-c	pyudev	configobj	pep8
amqp	happybase	os-refresh-c	PyYAML	coverage	proboscis
anyjson	httplib2	pam	qpidd-python	discover	psycopy2
argparse	iso8601	paramiko	requests	django-nose	pyflakes
Babel	Jinja2	passlib	Routes	docutils	pylint
boto	jsonpatch	Paste	rtplib-fb	feedparser	pysendfile
cffi	jsonpath-rw	PasteDeploy	simplejson	fixtures	pysqlite
cliff	jsonrpclib	pbr	six	flake8	python-ldap
coinor.pulp	jsonschema	pecan	sockjs-torna	hacking	python-subun
ddt	kazoo	pip	SQLAlchemy	hgtools	pymq
dogpile.cach	lesscpy	PrettyTable	sqlalchemy-m	hp3parclient	redis
diskimage-bu	kombu	psutil	stevedore	httpretty	selenium
Django	lockfile	pyasn1	suds	keyring	sphinx
django-boots	lxml	pycaddr	taskflow	mock	sphinxcontri
django_compr	msgpack-pyth	pycrypto	tripleo-imag	mox	sphinxcontri
django_opens	netaddr	pyghmi	warlock	mox3	sphinxcontri
dnspython	netifaces	pymongo	WebOb	MySQL-python	oslo.sphinx
eventlet	oauthlib	pyOpenSSL	websockify	nose	testreposito
extras	ordereddict	pyparsing	wheel	nose-exclude	testresource
falcon	oslo.config	pysnmp	wsgiref	nosehtmloutp	testscenario
Flask	oslo.messagi	pystache	WSME	nosexcover	testtools
futures	oslo.rootwra	python-memca	xattr	openstack-do	unittest2
gear	os-apply-con	pytz	cliff-tablib	openstack.no	WebTest

Project Packaging

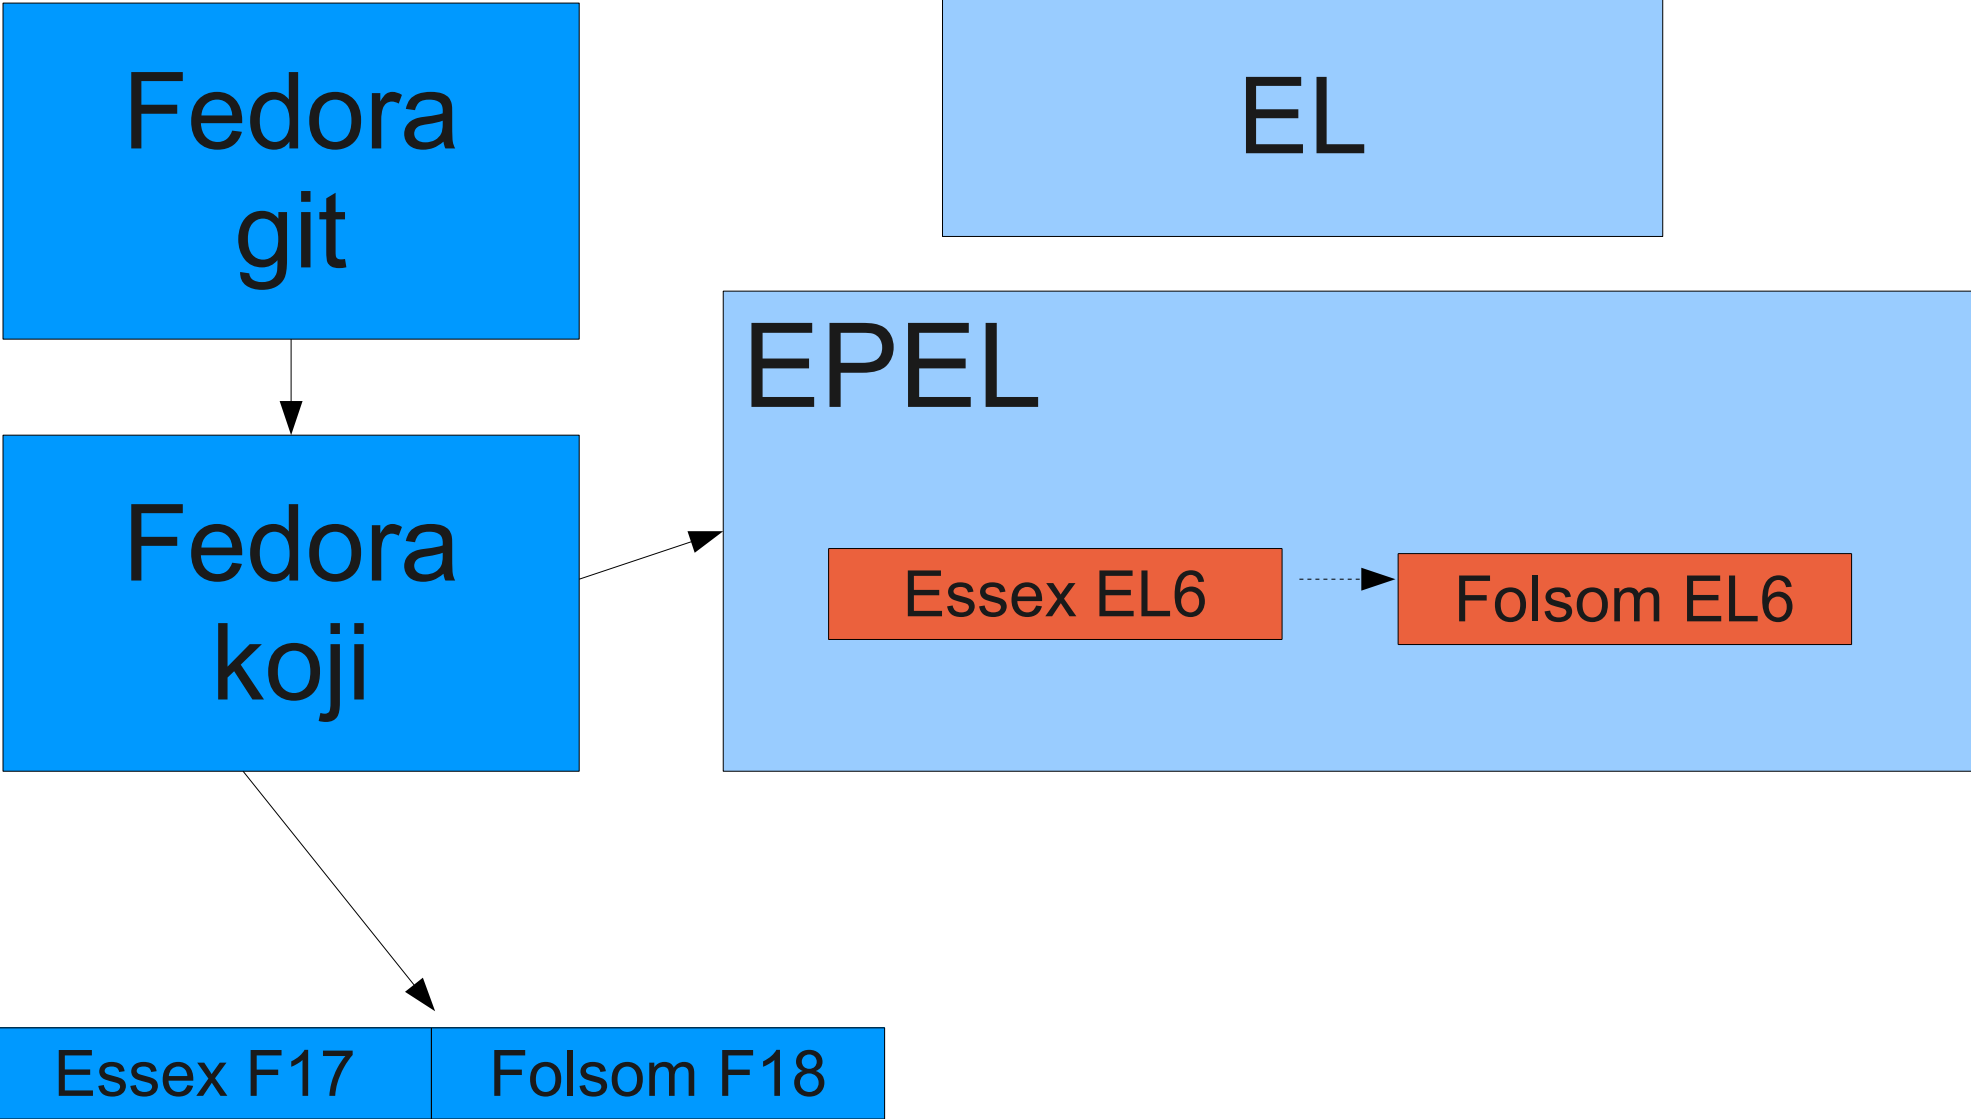
- Given the number of parts and options there is a large gain with consuming through distros
 - Often allows tweaks in the right place rather than workarounds in the wrong place
 - Leverages lots of logic and experience in testing, provisioning and upgrades
- BTW distro package deps are a good way to get overall position and leverage of a project

```
yum install rpmorphan graphviz  
rpmdep -dot openstack-deps.dot \  
  openstack-{nova,glance,cinder,ceilometer,dashboard,keystone,neutron}  
dot -Tsvg openstack-deps.dot -o openstack-deps.svg
```


OpenStack history on EL

- Essex ... Icehouse
- From a packaging viewpoint

OpenStack Essex/Folsom EL Repos



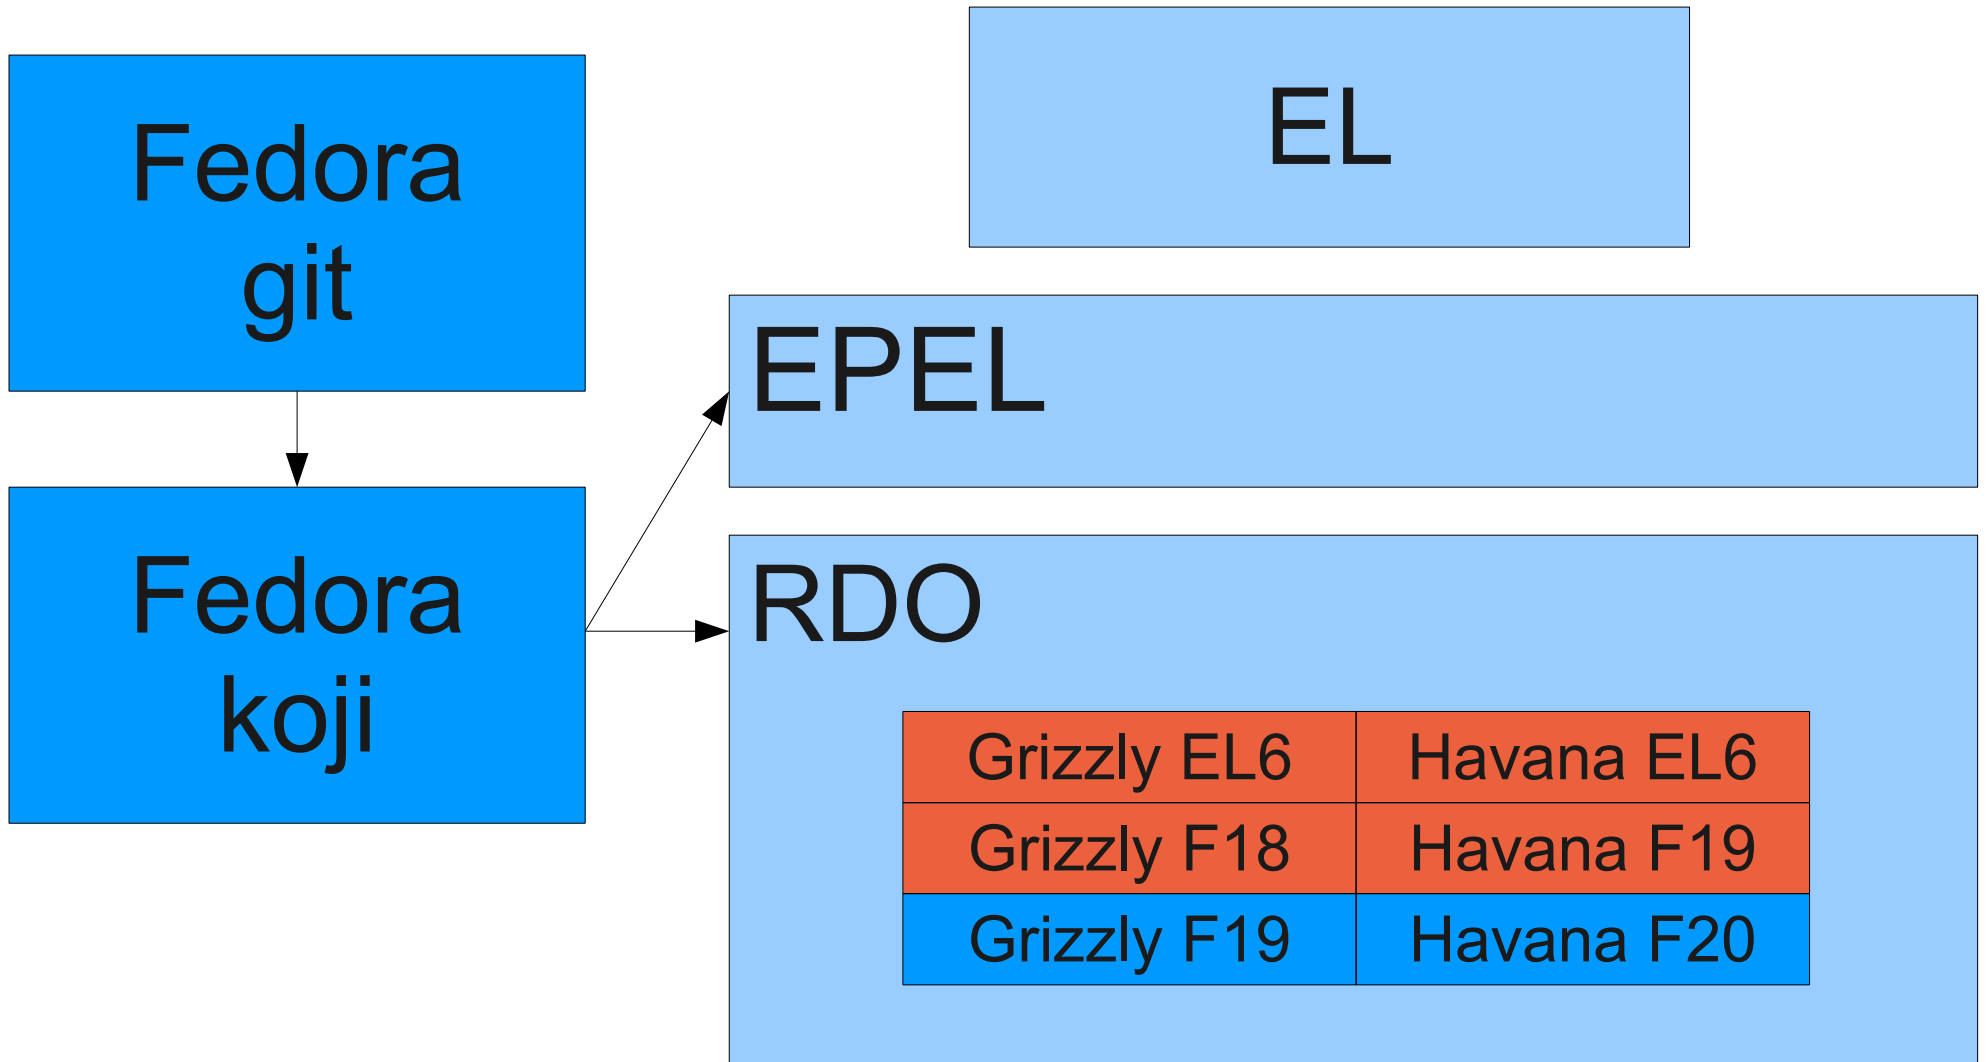
Issues with OpenStack in EPEL

- Upgrade is not seamless yet
 - Getting a lot better than Essex -> Folsom was but..
- All EL OpenStack users may not want to rebase
 - At least not at the same time
- Too restrictive for new dependencies
 - Parallel install packages awkward
 - OpenStack is large as we've seen so lots of potential overlap/conflict with general package set
 - Increasingly full compat is less of a requirement
 - With move to cloud and more ephemeral hosts

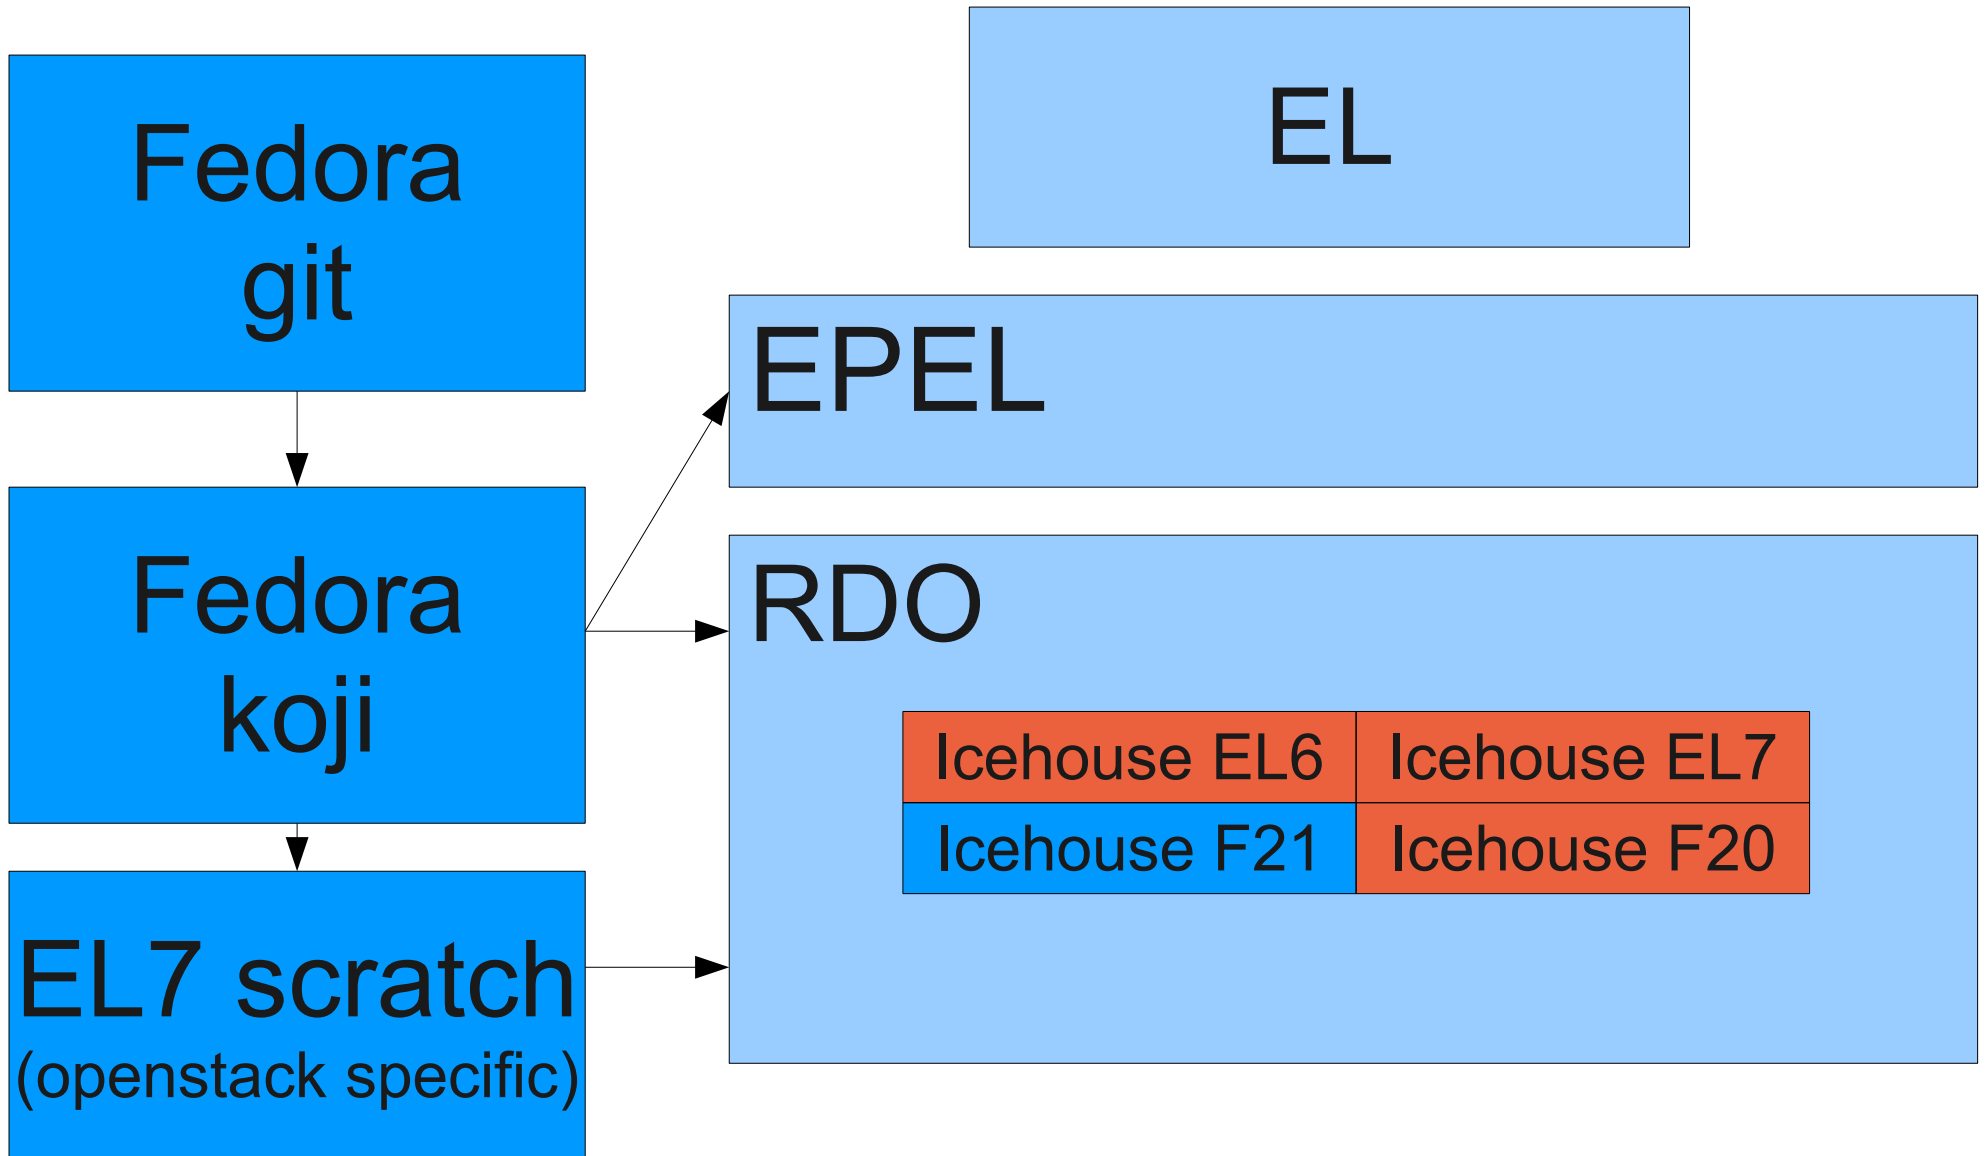
So enter RDO

- RDO is a **community** of people using and deploying OpenStack on Red Hat Enterprise Linux, Fedora and distributions derived from these (such as CentOS, Scientific Linux, ...)
- <http://openstack.redhat.com/>
 - 3 step install process for CentOS etc.
- Part of that are separate more flexible repos
 - Essentially leveraging the fact that OpenStack installs are almost always dedicated hosts
 - Hence global compat is desired but not required
 - Reference thirdparty repos like puppetlabs.org

OpenStack Grizzly/Havana EL Repos



OpenStack Icehouse EL Repos



Changes for Icehouse EL7

- More detachment from EPEL7
 - No OpenStack specific dependencies in EPEL7
- OpenStack **build** dependencies
 - Added to EPEL7 to allow all builds in Koji
 - There are some openstack **related** but not openstack **specific**
 - python-pbr
 - python-testtools
 - python-pip
 - python-mimeparse
 - python-extras
 - [python-testtools]
 - python-oslo-sphinx
 - crudini
- OpenStack specific packages
 - Koji scratch builds for now (copied to RDO)

Changes for Icehouse EL7

- OpenStack specific runtime dependencies
 - Koji scratch builds for now (copied to RDO)
 - python-oslo-{config,rootwrap,messaging}
- New non OpenStack specific runtime deps
 - Added to EPEL7
- Updated non OpenStack specific runtime deps
 - python-six-1.4.1-1.el7 (newer than EL7)
 - Can't go in EPEL7
 - Koji scratch builds for now
 - Candidate for an “update repo” for progression to base
 - Similar to updated kernel used with 6.4 which migrated to 6.5 base

Possible CentOS structure to help

- git repos and buildroots to handle the koji scratch build cases above
- Shared repos, like “cloud sig repo”
 - Instance
 - Cloud-init
 - Infrastructure
 - Shared puppet
 - Overlap with other sigs (java, python?)
- I can see the need for packages that are not compatible with all other packages in the major EL release. (CentOS next?)
- More usually I can see the need for packages that are updates for existing packages, which may be candidates for future updates in EPEL or even an EL minor release (CentOS updates?)

Possible shared CentOS repos

EL

EPEL

CentOS
Updates

CentOS
Next

OpenStack

OpenShift

CloudStack

Notes on CentOS repos

- Packages would aim to move up to increase sharing and remove duplication
- Updates for existing packages would provide useful feedback for update candidates in base
- New packages are only new once, and hence really have the same constraints as “Updates”
- Tracking multiple sigs using new packages would increase pressure for those packages to move up
- Should aim to minimize levels to avoid admin overhead and to ease the flow of packages up