


Managing PostgreSQL with Puppet



CCBYSA-NC <http://www.flickr.com/photos/urgetopunt/7876067126/>

Steve Singer
ssinger@ca.afilias.info

Who is Steve?



- Slony maintainer
- Database Developer @ Afilias (Toronto)



Afilias™



DNS
Lookup




Registrars



100+ of Databases



Technologies



DevOps



Deployment Automation




Goals of Automation: Database servers 'look' the same



Consistency between QA, Staging, Production



Reduce Manual Commands



Puppet

http://puppetlabs.com

File Edit View History Bookmarks Tools Help

Puppet Node Manager

https://puppet-ent/?ticket=ST-1366990345rD6A926C7F06259DAC9

Getting Started Finance Latest Headlines twitter

puppet enterprise console node requests (0) ssinger Help

Nodes Groups Classes Reports File Search Inventory Search Compliance Live Management

Background Tasks

- All systems go

Nodes

- 309 Unresponsive
- 4 Failed
- 0 Pending
- 3 Changed
- 1 Unchanged
- 7 Unreported
- 324 All**

Add node Radiator View

Group

- DB Servers (3)
- default (324)

Add group

Class

Daily run status

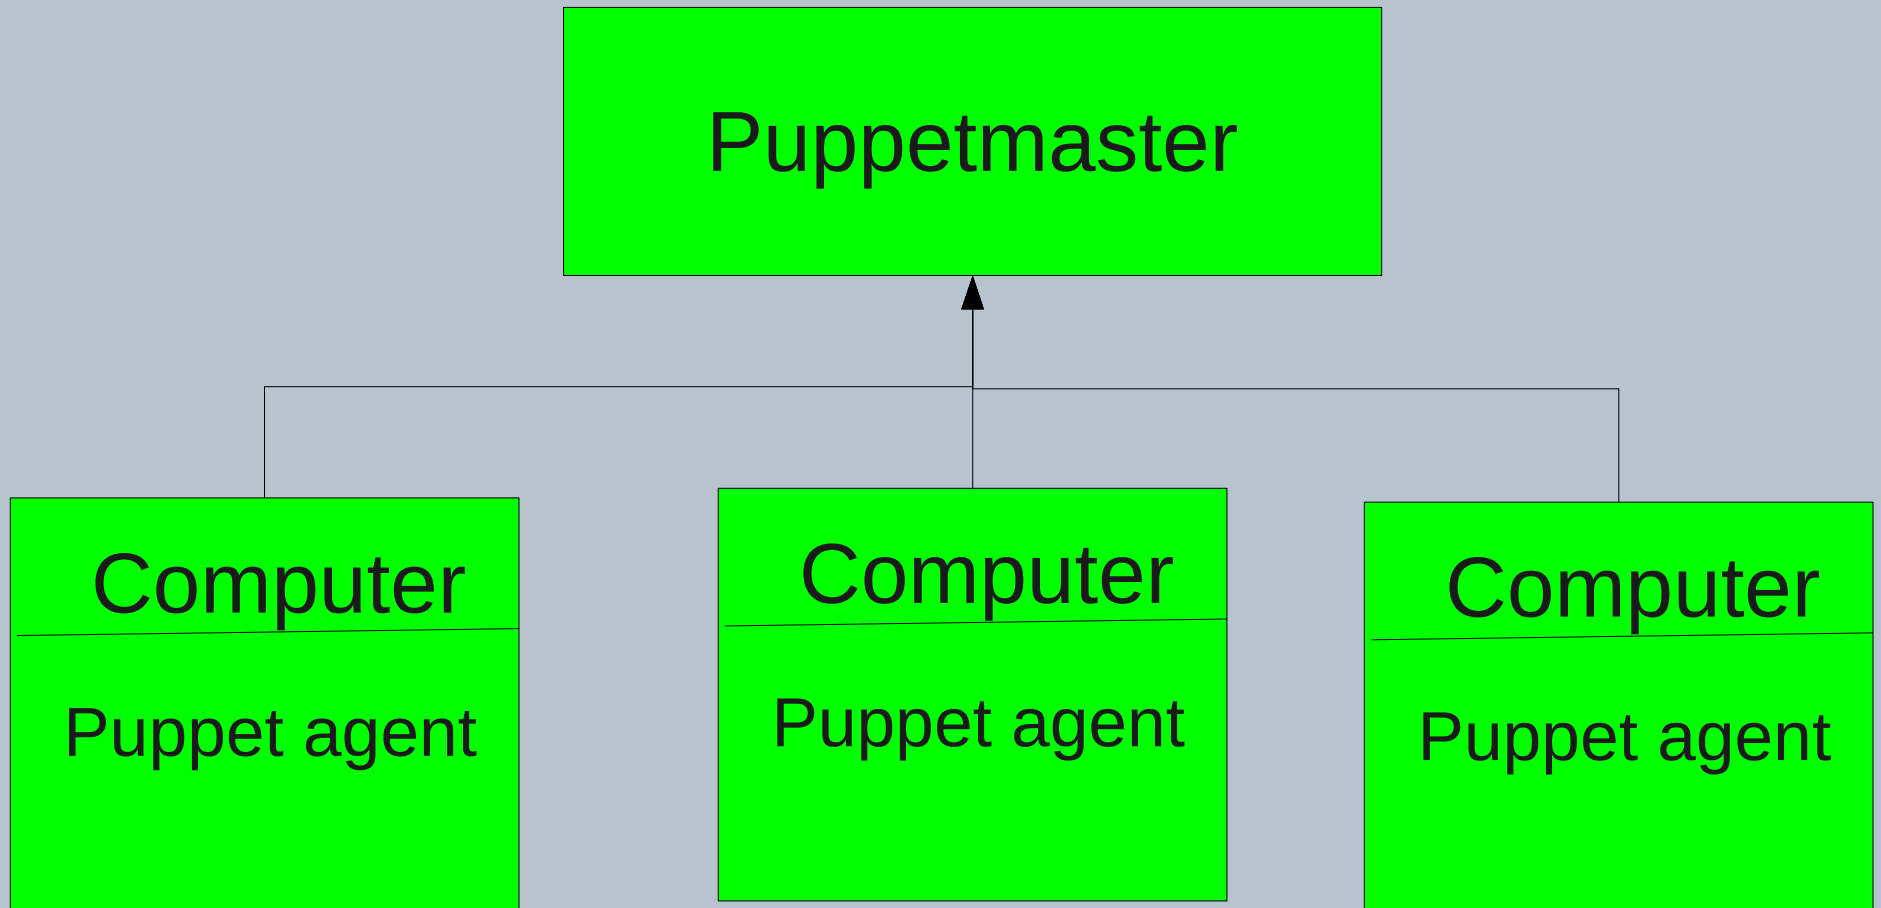
Number and status of runs during the last 30 days:

All Unresponsive Failed Pending Changed Unchanged

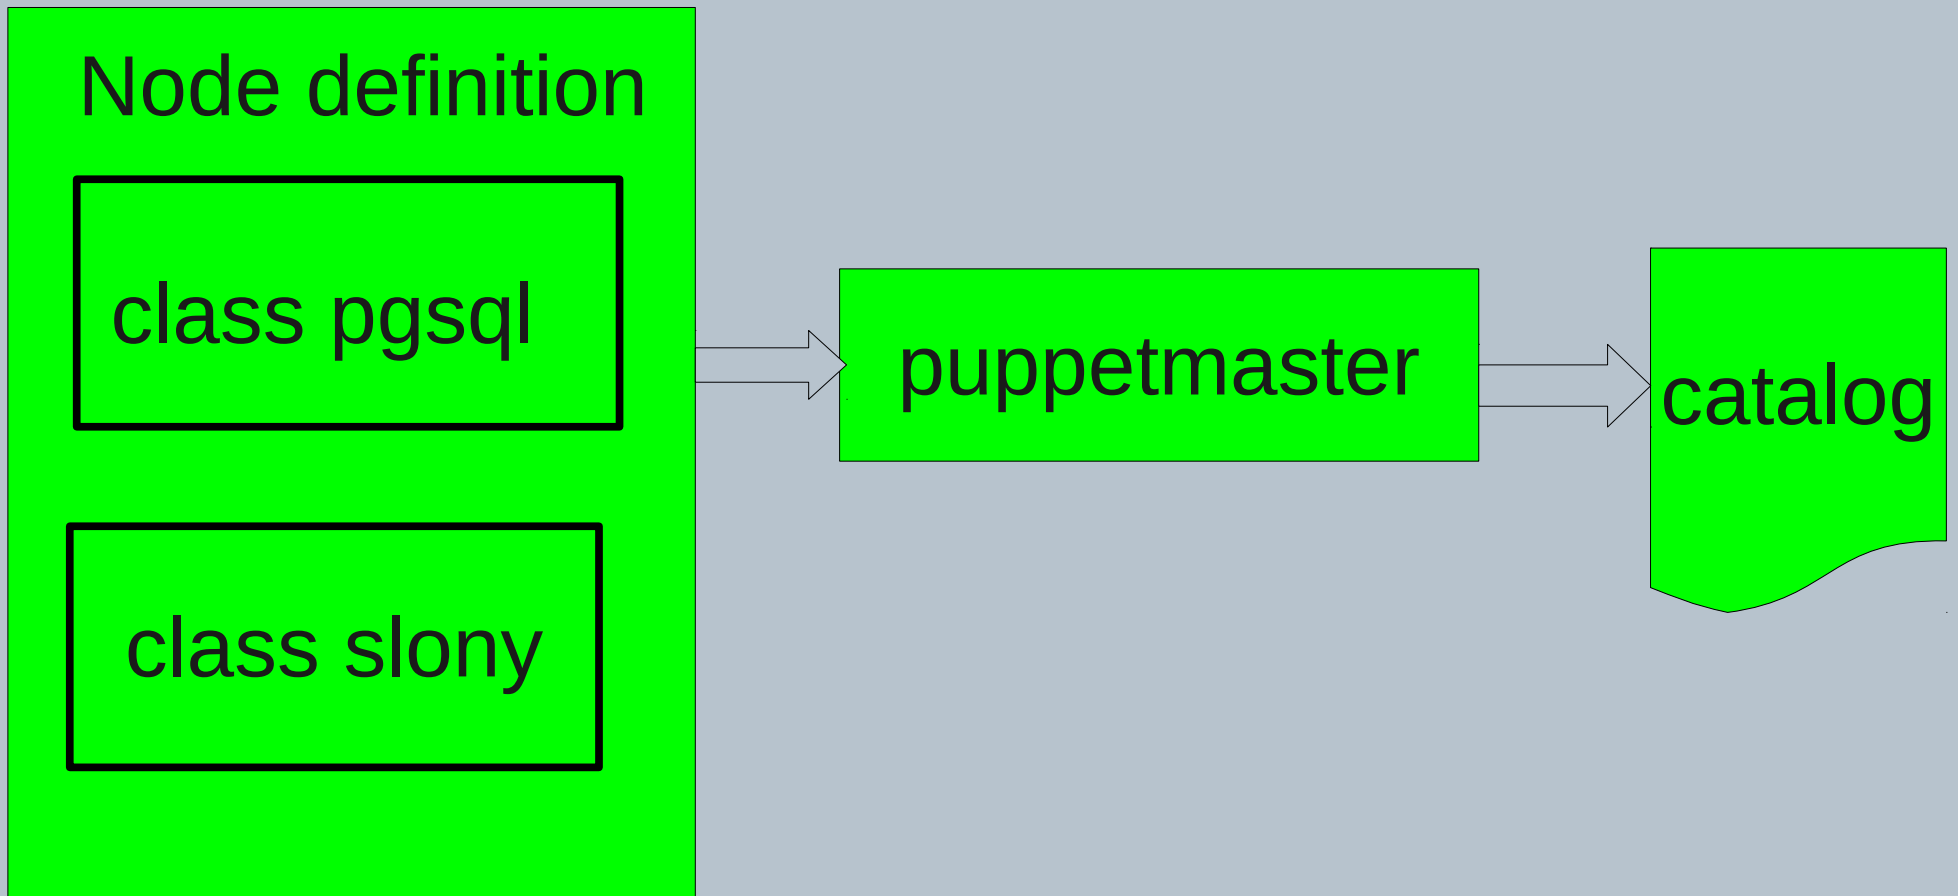
Export nodes as CSV

Node	Latest report	Resources				
		Total	Failed	Pending	Changed	Unchanged
Total		77130	91	166	2286	74650
oteepp3.on1.afllias-ops.info	2013-04-26 15:30 UTC	449	0	0	0	449
dap2.tx1.afllias-ops.info	2013-04-26 15:26 UTC	1044	2	0	141	901
dap1.on1.afllias-ops.info	2013-04-26 15:22 UTC	1030	0	0	138	892
zenosrm1.tor.afllias-int.info	2013-04-26 15:21 UTC	0	0	0	0	0
dap1.tx1.afllias-ops.info	2013-04-26 15:17 UTC	798	0	0	92	706
dap3.tx1.afllias-ops.info	2013-04-26 15:07 UTC	721	0	0	85	636
dap2.on1.afllias-ops.info	2013-04-26 15:06 UTC	1075	4	0	214	857

Puppet Architecture



Puppet Terminology



Puppet Class

```
class postgresql {  
  package { 'postgresql-server' :  
    ensure => installed,  
    provider => 'yum'  
  }  
}
```

resource

Resource Type

```
graph TD  
  resource --> package  
  RT[Resource Type] --> package
```

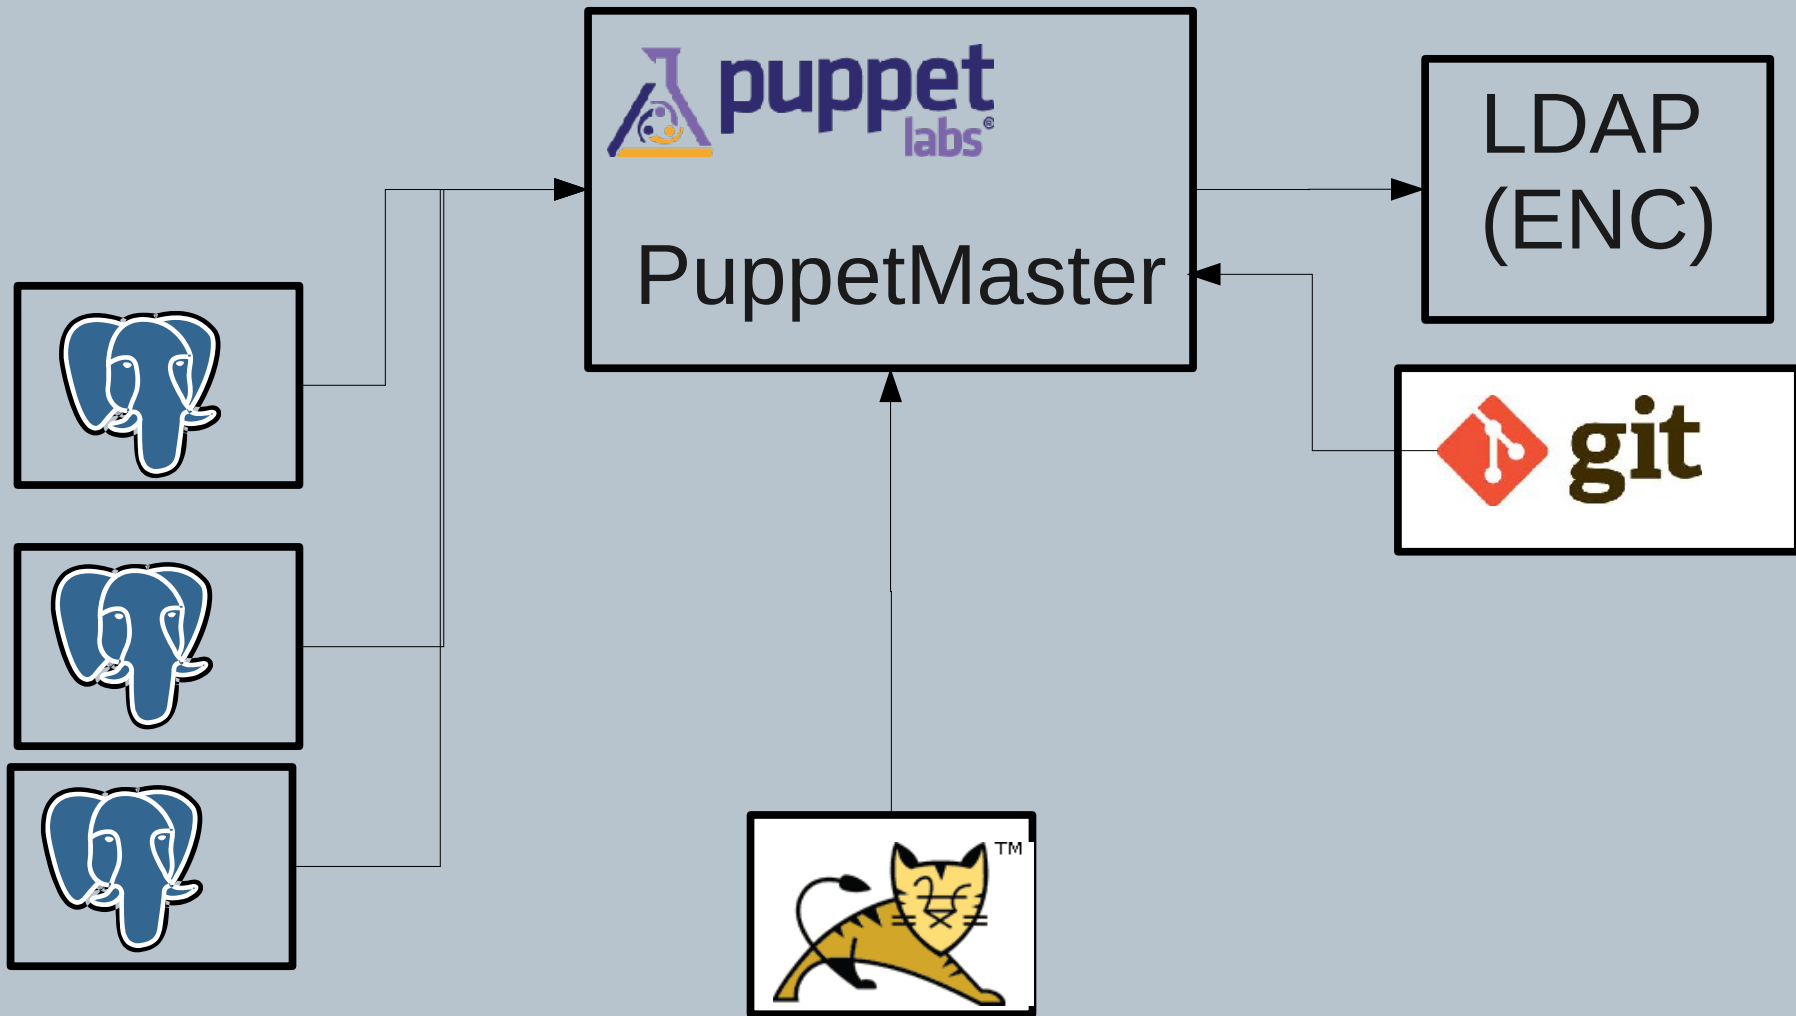

Declarative

```
file { '/tmp/hello_world.txt' :  
  content => 'Welcome to Chicago',  
  owner => 'pgopen',  
  mode => 0644  
}
```



```
/tmp/hello_world.txt  
Welcome to Chicago
```

Puppet At Afilias



Goals for Databases

- Deploy Binaries
- Manage `pg_hba.conf` and `postgresql.conf`
- Manage postgres unix users
- Handle init db
- Manage Crontabs
- Manage slons and slonik preambles

PostgreSQL Module

<https://github.com/puppetlabs/puppet-postgresql>

- Pre-existing module for managing postgresql
- Manages installing, start/stopping postgresql
- We don't actually use it

pgsql_cluster resource

```
define pgsql_cluster ($clustername,$port,$datadir,  
    $conf_template,$listen_address,$pguser,  
    $pghba_template,$service_name,$pguser_password,$encoding)  
{  
    ...  
}
```

- One instance per database \$datadir
- Performs initdb
- Deploys postgresql.conf
- Deploys pg_hba.conf
- Creates standard database users

Templates

postgresql.conf.erb

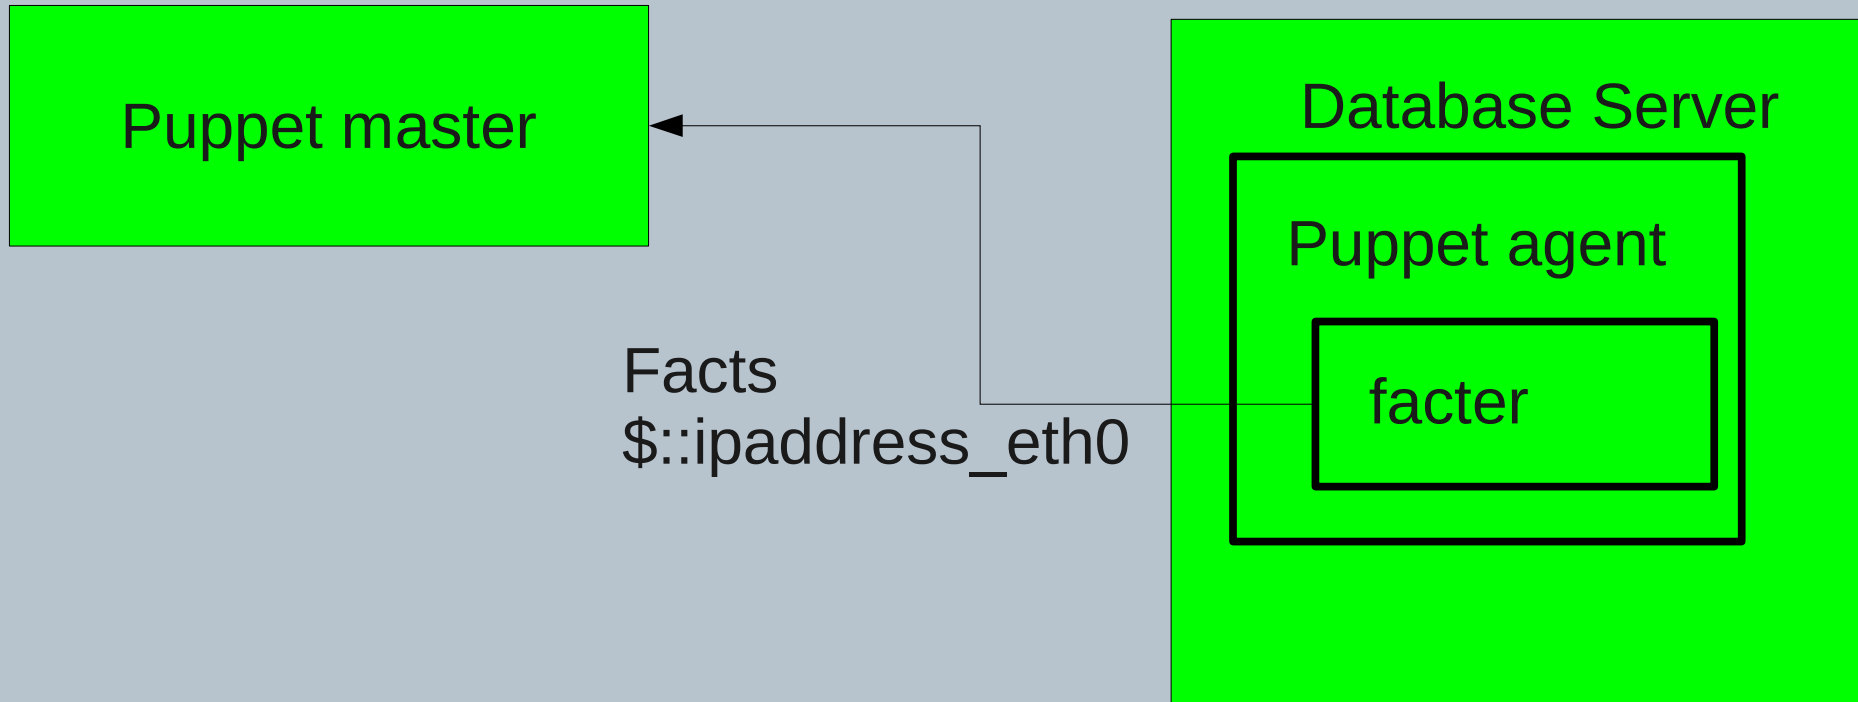
```
#-----  
# RESOURCE USAGE (except WAL)  
#-----  
  
# - Memory -  
  
shared_buffers = <%= shared_buffers %>    # min 128kB  
# (change requires restart)
```


postgresql.conf

```
$shared_buffers= '300MB'  
$conf_template='dba/postgresql.91.conf.erb'  
  
file { "${pg_datadir}/postgresql.conf" :  
    owner    => "$pguser",  
    group    => postgres,  
    mode     => 644,  
    require => Exec["${name}_initdb"],  
    content => template("$conf_template"),  
}
```

listen_addresses and facter

```
listen_address= '<%= scope.lookupvar('::ipaddress_eth0') %>'
```



Hiera 101

Separate Code From Data



http://commons.wikimedia.org/wiki/File:Lower_Antelope_Canyon_478.jpg
By Uploaded by Meckimac [GFDL (<http://www.gnu.org/copyleft/fdl.html>) or CC-BY-SA-3.0 (<http://creativecommons.org/licenses/by-sa/3.0/>)],
via Wikimedia Commons

Hiera 101: Lookup Values

```
class db_server {  
    $application='postgres'  
    $product='billing'  
    $port = hiera('port')  
    .  
    .  
}
```


Hiera 101: Define a Hierarchy

```
--  
:hierarchy:  
-domains/%{domain}/%{fqdn}/%{application}  
-domains/%{domain}/%{application}/  
-environment_type/%{environment_type}/%{application}  
-complexes/%{product}/%{application}
```

Hiera 101: Hierarchy

domains/tor.afili-as-int.info/db1.tor.afili-as-int.info

postgres.yaml

db2.tor.afili-as-int.info

postgres.yaml

environment_type
/QA

postgres.yaml

PRODUCTION

postgres.yaml

products

billing

postgres.yaml

wiki

postgres.yaml

Hiera 101: postgres.yaml

```
port : 5432
```

```
postgres_pgversion : '9.2.4'
```

```
postgres_unixuser : postgres
```

```
postgres_db_list :
```

```
  pgopen : ssinger
```

```
  pgeu : jwieck
```

conf template server type

postgresql.90.small.conf.erb

postgresql.91.small.conf.erb

postgresql.92.small.conf.erb

postgresql.90.oltp.conf.erb

postgresql.91.oltp.conf.erb

postgresql.92.oltp.conf.erb

postgresql.90.batch.conf.erb

postgresql.91.batch.conf.erb

postgresql.92.batch.conf.erb

Switch Based Config

postgresql.90.conf.erb

```
shared_buffers = <%= shared_buffers %>  
max_connections = <%= max_connections %>
```

postgresql.91.conf.erb

```
shared_buffers = <%= shared_buffers %>  
max_connections = <%= max_connections %>
```

postgresql.92.conf.erb

```
shared_buffers = <%= shared_buffers %>  
max_connections = <%= max_connections %>
```

Crontabs

```
cron { "backup_${name}" :  
    command => " /opt/dbs/scripts/backup.sh $clustername",  
    user => 'postgres',  
    hour => '0',  
    minute => '0',  
    ensure => 'present'  
}
```


Slony

```
cluster name = mydbcluster  
node 9181 admin conninfo = 'host=db1.on1 user=slony'  
node 9182 admin conninfo = 'host=db2.on1 user=slony'  
  
set add table ( id =1 , fully qualified name = 'public.foo' );
```

Slony Settings

```
# YAML for the slons running on a particular server
postgres_slons:
- 9181
```

```
# YAML map of connection data to all slons
postgres_slon_nodes:
  '9181':
    'host': 'db1.on1.afilias-int.info'
    'port' : '5432'
  '9192':
    'host': 'db2.tx1.afilias-int.info'
    'port' : '5432'
```

Slony: Slonik preambles

```
/opt/configs/$clustername/slonik.preamble :  
slonik_preamble.erb
```

```
cluster name=<%= clustername %>;  
# A slonik preamble that defines  
# the 'node xxx admin conninfo...' data  
# and an associated DEFINE for each  
# node in the cluster  
#  
<%= slon_conf.each do | node_id, values | %>  
<%= port=values['port']  
  if port==nil or port=="  
    port=complex_port  
  end %>  
node <%= node_id %> admin conninfo='host=<%= values['host'] %> .....  
define CONNINFO_<%= node_id %> 'host=<%= values['host'] %> .....  
<%= end %>
```

Removing Things

```
class db_node {  
  .  
  .  
}  
  
class db_node::remove {  
  file { "$datadir",  
        ensure => 'absent',  
        force => true,  
        Backup => false  
  }  
  .  
  .  
}
```

pg_basebackup

```
if($datadir_restore_command) {
  exec { "${name}_initdb" :
    command => $datadir_restore_command,
    unless => "/usr/bin/test -d ${pg_datadir}/base"
    ...
  }
}
else {
  exec "${name}_initdb" :
    command => /usr/bin/initdb -D ${pg_datadir}
    unless => "/usr/bin/test -d ${pg_datadir}/base"
  }
}
```

Working with others



Questions?



<http://puppetlabs.com>
ssinger@ca.afilias.info