



# CHEF IN THE CLOUD AND ON THE GROUND

Michael T. Nygard  
*Relevance*

[michael.nygard@thinkrelevance.com](mailto:michael.nygard@thinkrelevance.com)

[@mtnygard](https://twitter.com/mtnygard)

INTERNATIONAL  
SOFTWARE DEVELOPMENT  
CONFERENCE

[gotocon.com](http://gotocon.com)

# Infrastructure As Code

# Infrastructure As Code

## Chef

# Infrastructure As Code

## Chef

# Development Models

# Infrastructure As Code

## Chef

### Development Models

### Deployment Models

# Infrastructure As Code

# Infrastructure As Code

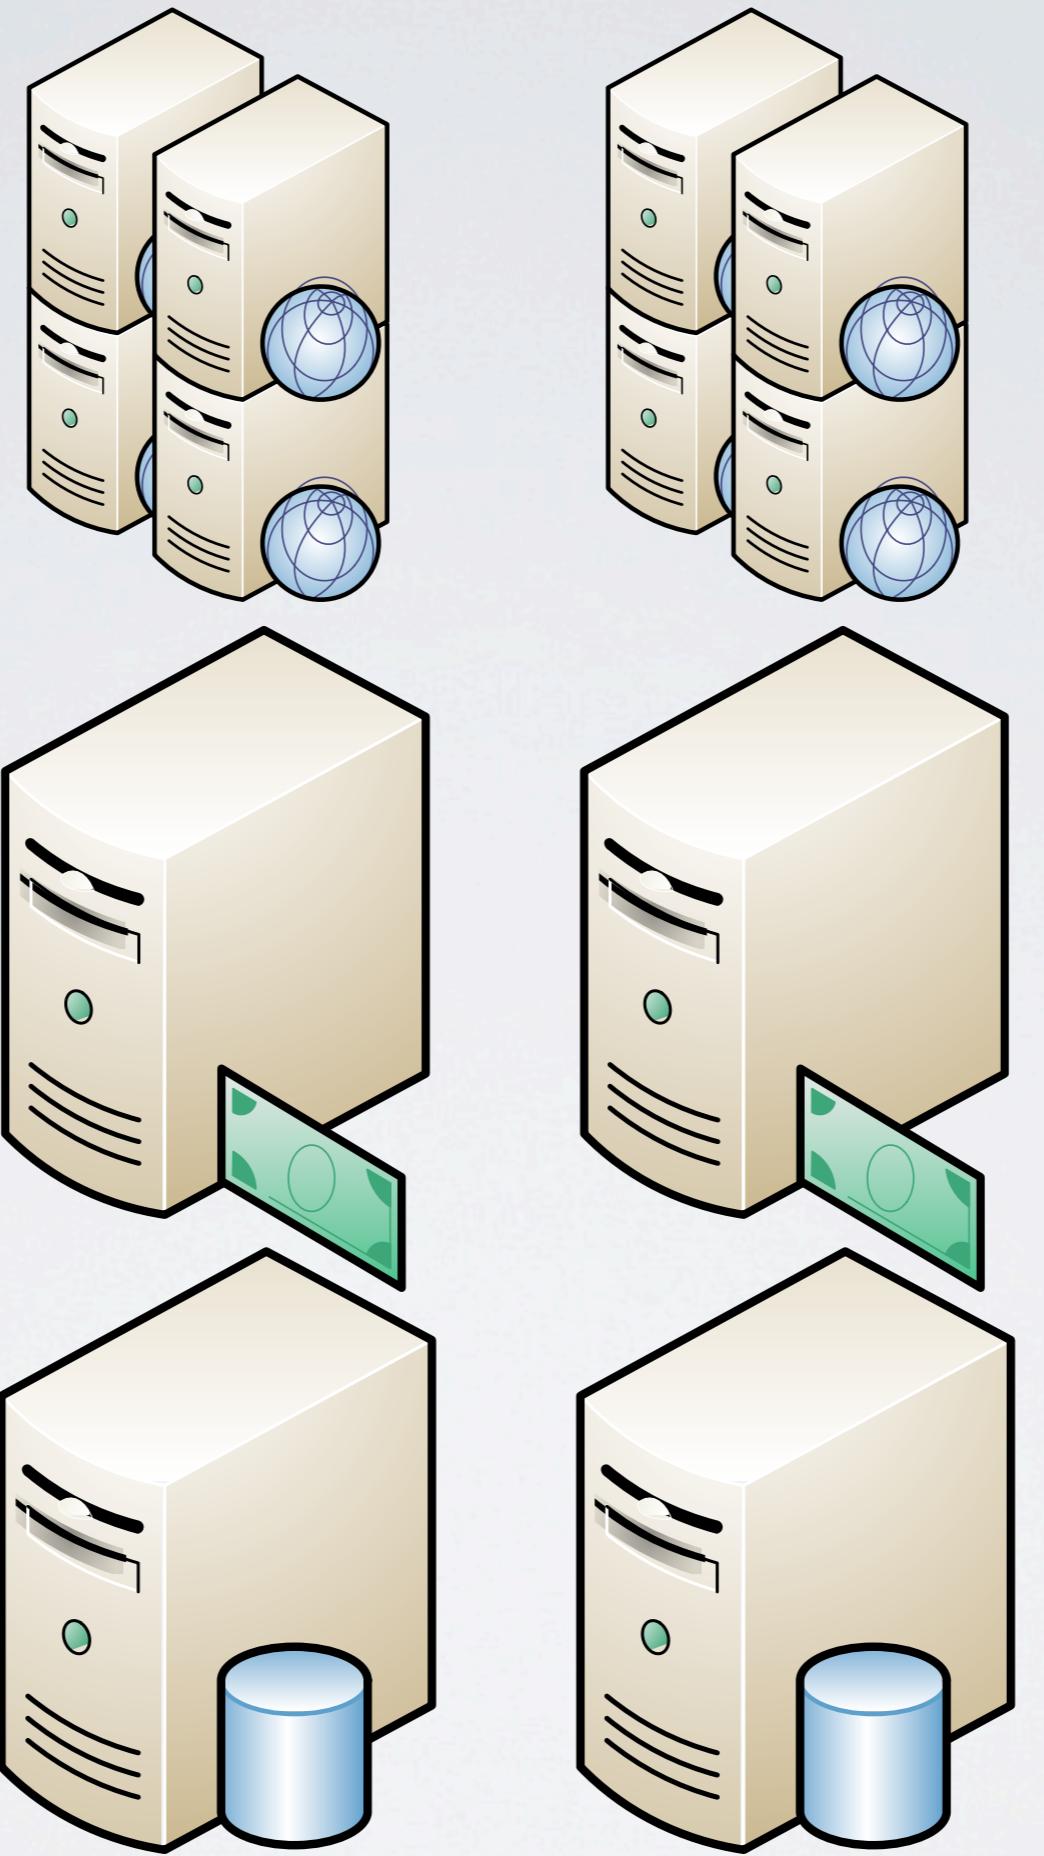
Problems beget solutions.  
Solutions become the next problems.



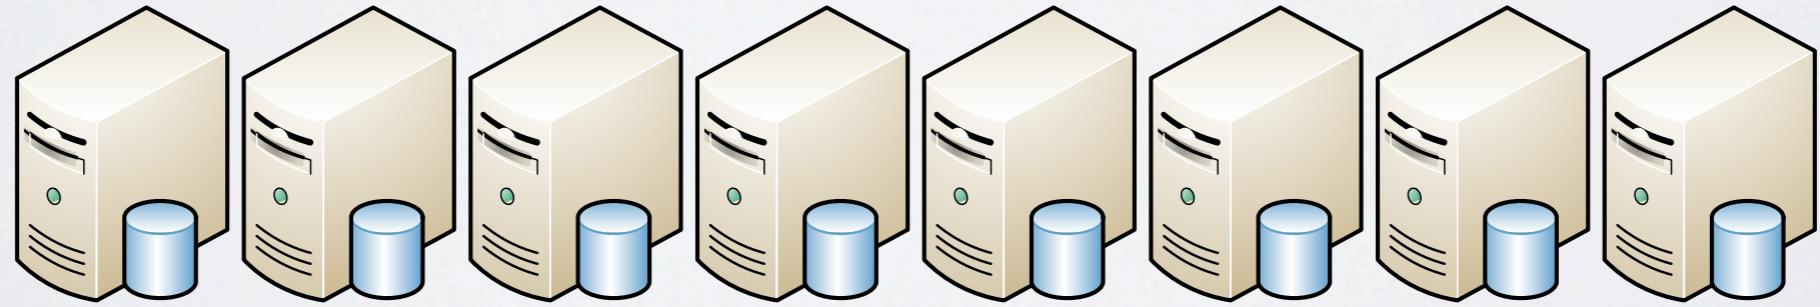
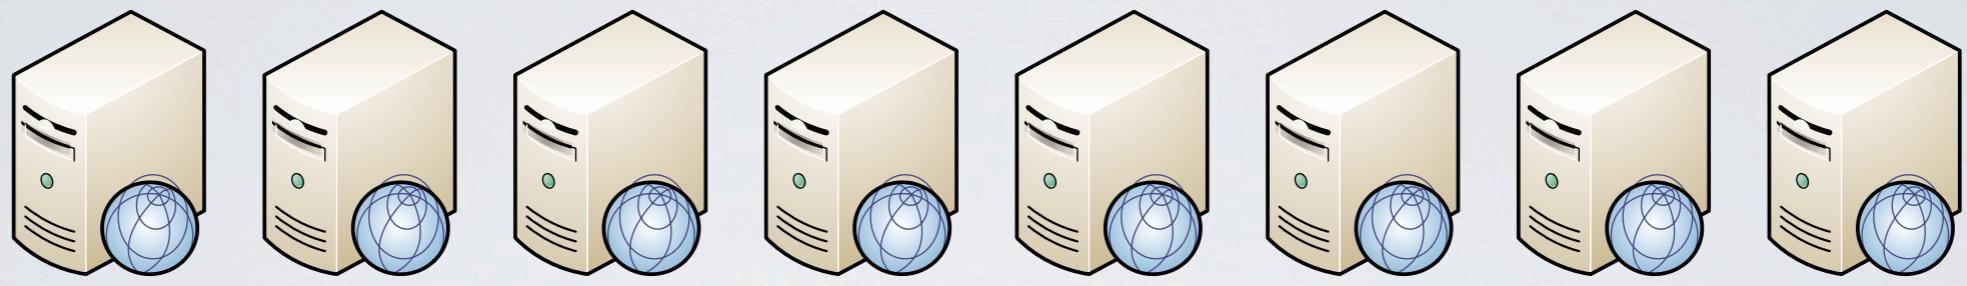












What if we could version control  
our servers and network, the  
way we do with code ?

# Infrastructure as Code

# Infrastructure as Code

# Infrastructure as Code

- Desired state

# Infrastructure as Code

- Desired state
- Transformations to achieve it

# Infrastructure as Code

- Desired state
- Transformations to achieve it
- As text files

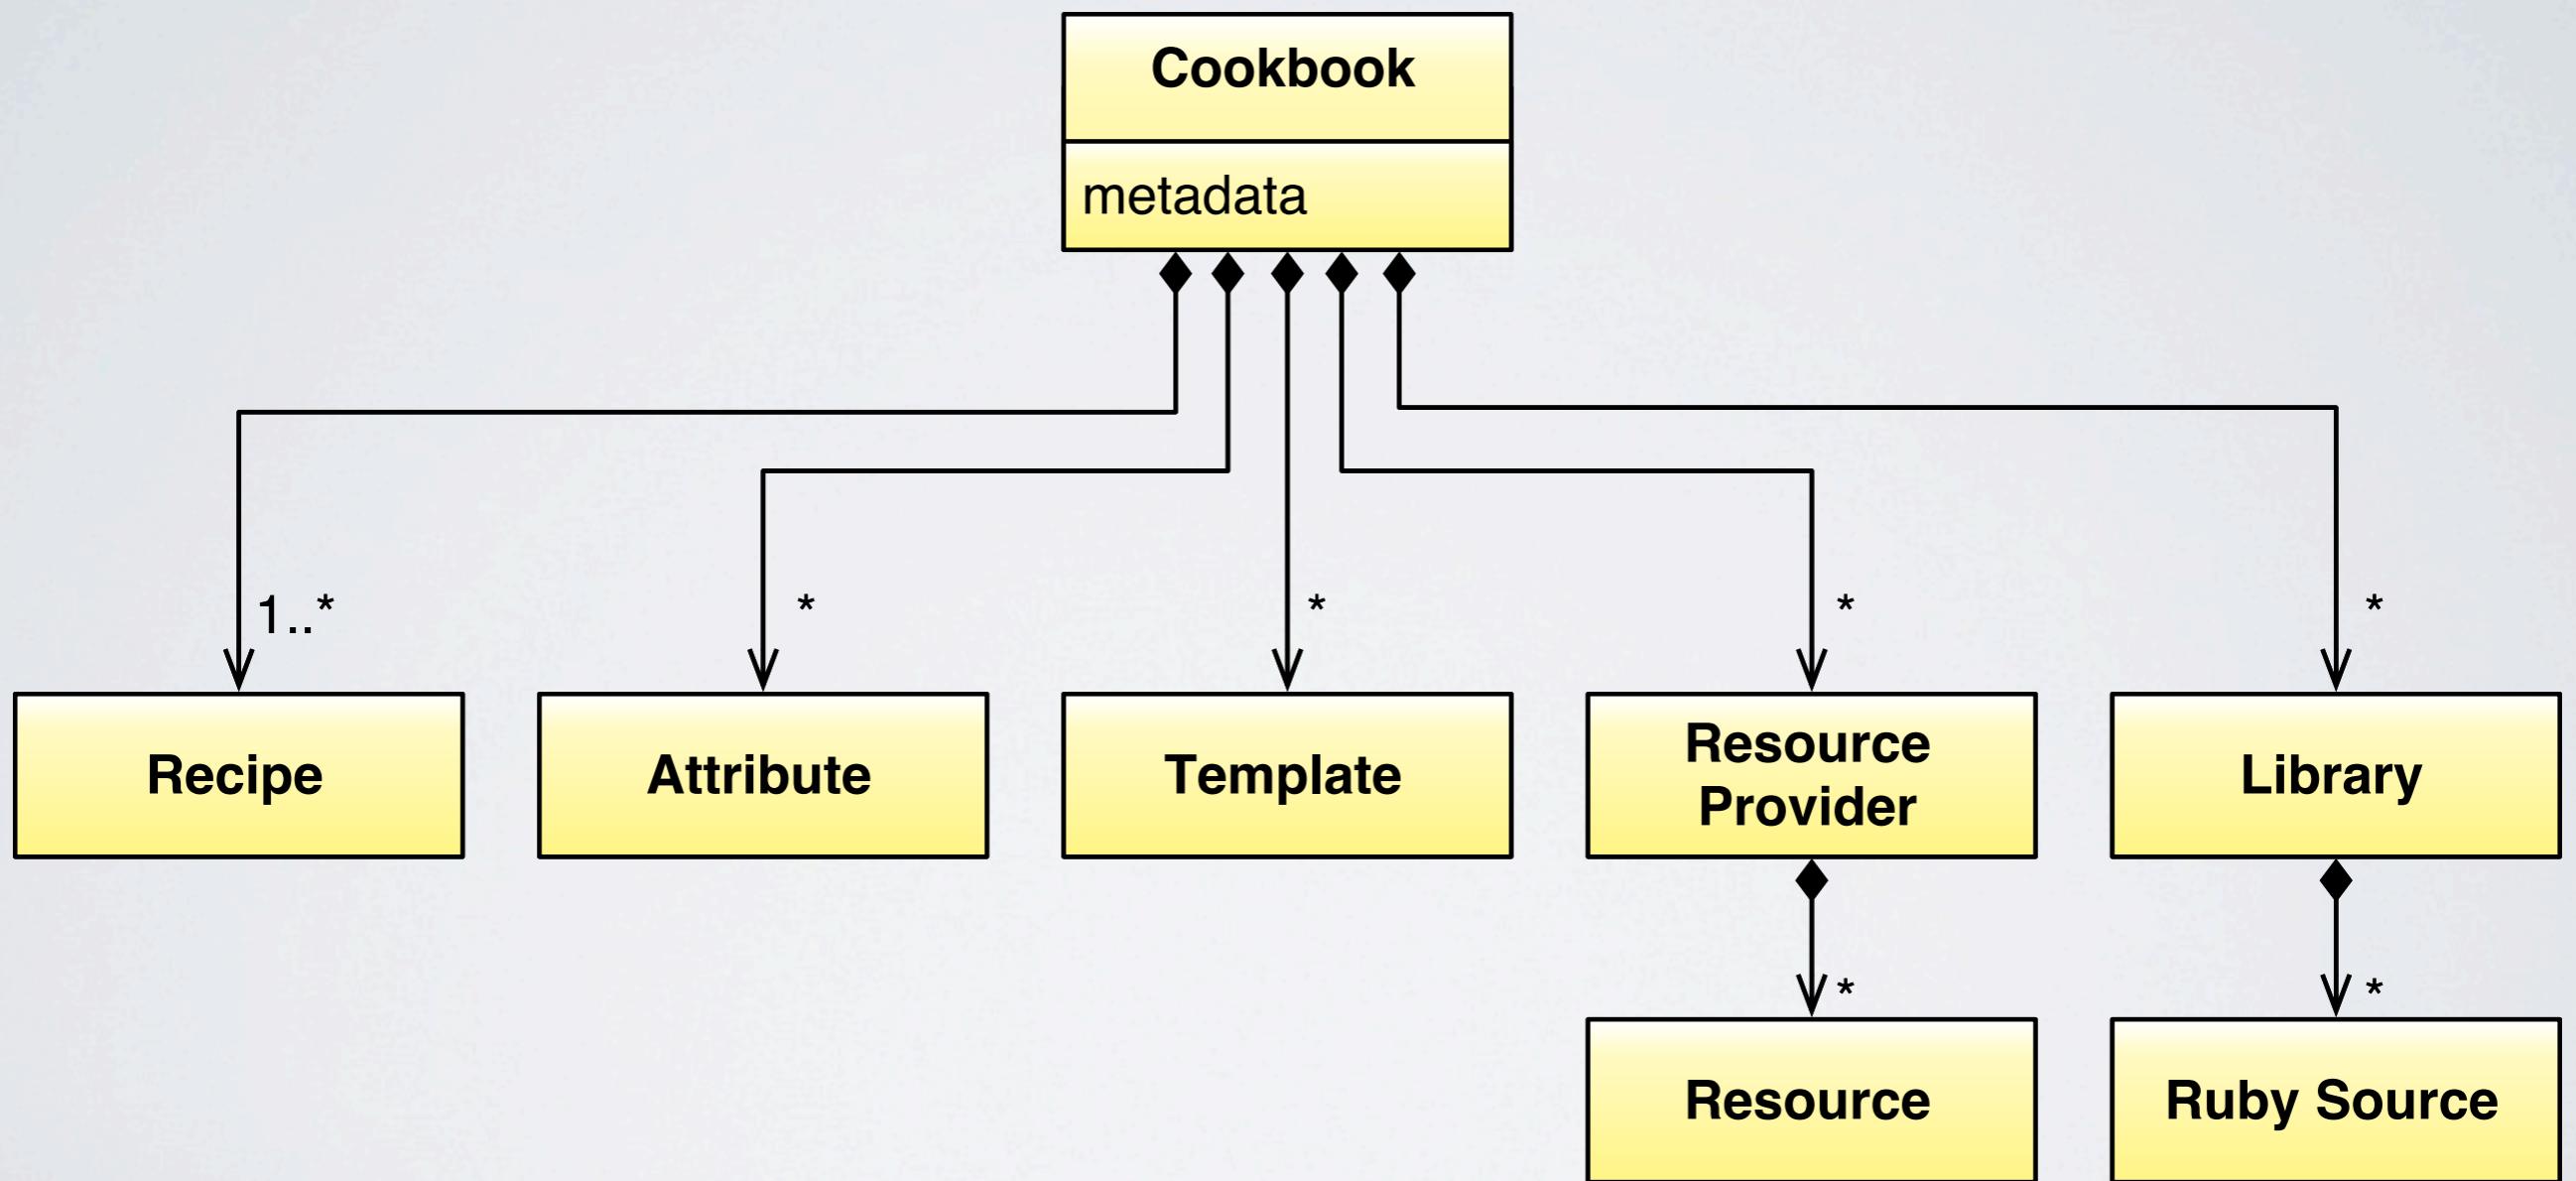
# Infrastructure as Code

- Desired state
- Transformations to achieve it
- As text files
- In version control

# Infrastructure as Code

- Desired state
- Transformations to achieve it
- As text files
- In version control
- Testable and reproducible

# Chef Basics



# Chef Basics

# Chef Basics

Cookbook

Container for semantically related recipes and their support files.

# Chef Basics

Cookbook	Container for semantically related recipes and their support files.
Recipe	Resource definitions, written as Ruby code

# Chef Basics

Cookbook	Container for semantically related recipes and their support files.
Recipe	Resource definitions, written as Ruby code
Attribute	Parameter used in recipes, can be supplied by cookbook defaults, the node, or the role

# Chef Basics

Cookbook	Container for semantically related recipes and their support files.
Recipe	Resource definitions, written as Ruby code
Attribute	Parameter used in recipes, can be supplied by cookbook defaults, the node, or the role
Template	ERB file that gets expanded onto the target host, used by resources in recipes

# Chef Basics

Cookbook	Container for semantically related recipes and their support files.
Recipe	Resource definitions, written as Ruby code
Attribute	Parameter used in recipes, can be supplied by cookbook defaults, the node, or the role
Template	ERB file that gets expanded onto the target host, used by resources in recipes
Library	Extra Ruby code, can create classes, resources, implement actions, etc.

# Chef Basics

Cookbook	Container for semantically related recipes and their support files.
Recipe	Resource definitions, written as Ruby code
Attribute	Parameter used in recipes, can be supplied by cookbook defaults, the node, or the role
Template	ERB file that gets expanded onto the target host, used by resources in recipes
Library	Extra Ruby code, can create classes, resources, implement actions, etc.
Resource Provider	Lightweight Ruby DSL for creating new resources.

# Cookbook: Git

 git	
 metadata.json	
 metadata.rb	
 README.rdoc	
 recipes	
 default.rb	
 server.rb	
 templates	
 default	
 sv-git-daemon-log-run.erb	
 sv-git-daemon-run.erb	



## metadata.rb

```
maintainer          "Opscode, Inc."
maintainer_email    "cookbooks@opscode.com"
license             "Apache 2.0"
description         "Installs git and/or set..."
long_description   IO.read(File.join(
                           File.dirname(__FILE__),
                           'README.rdoc'))
version            "0.9.0"
recipe              "git", "Installs git"
                   "git::server",
                   "Sets up a runit_service for git daemon"
```

cont...



## metadata.rb

...

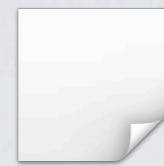
```
%w{ ubuntu debian arch}.each do |os|  
  supports os  
end
```

```
%w{ runit }.each do |cb|  
  depends cb  
end
```



## recipes/default.rb

```
case node[:platform]
when "debian", "ubuntu"
  package "git-core"
else
  package "git"
end
```



## recipes/server.rb

```
include_recipe "git"

directory "/srv/git" do
  owner "root"
  group "root"
  mode 0755
end
```

cont...



## recipes/server.rb

...

```
case node[:platform]
when "debian", "ubuntu"
  include_recipe "runit"
  runit_service "git-daemon"
else
  log "Platform requires setting up ..."
  log "Hint: /usr/bin/git daemon --export..."
end
```

# Cookbook: MySQL

 mysql		
 attributes	 recipes	 templates
 server.rb	 client.rb	 default
 libraries	 default.rb	 debian.cnf.erb
 database.rb	 server_ec2.rb	 grants.sql.erb
 helpers.rb	 server.rb	 my.cnf.erb
 metadata.json	 resources	 mysql-server.seed.erb
 metadata.rb	 database.rb	 port_mysql.erb
 providers		
 database.rb		
 README.md		



## templates/default/debian.cnf.erb

```
[client]
host      = localhost
user      = debian-sys-maint
password  = <%= node['mysql']['server_debian_password'] %>
socket    = <%= node['mysql']['socket'] %>

[mysql_upgrade]
host      = localhost
user      = debian-sys-maint
password  = <%= node['mysql']['server_debian_password'] %>
socket    = <%= node['mysql']['socket'] %>
basedir   = /usr
```



## resources/database.rb

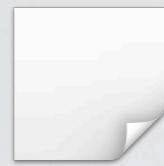
```
actions :flush_tables_with_read_lock,  
       :unflush_tables, :create_db, :query  
  
attribute :host, :kind_of => String  
attribute :username, :kind_of => String  
attribute :password, :kind_of => String  
attribute :database, :kind_of => String  
attribute :sql, :kind_of => String  
attribute :exists, :default => false
```

# Using the Custom Resource

```
# In cookbooks/stratego/recipes/default.rb

include_recipe "mysql::client"

mysql_database "reporting database" do
  host "db01prod"
  username "root"
  password node[:mysql][:server_root_password]
  database "reporting_production"
  action :create_db
end
```



## providers/database.rb

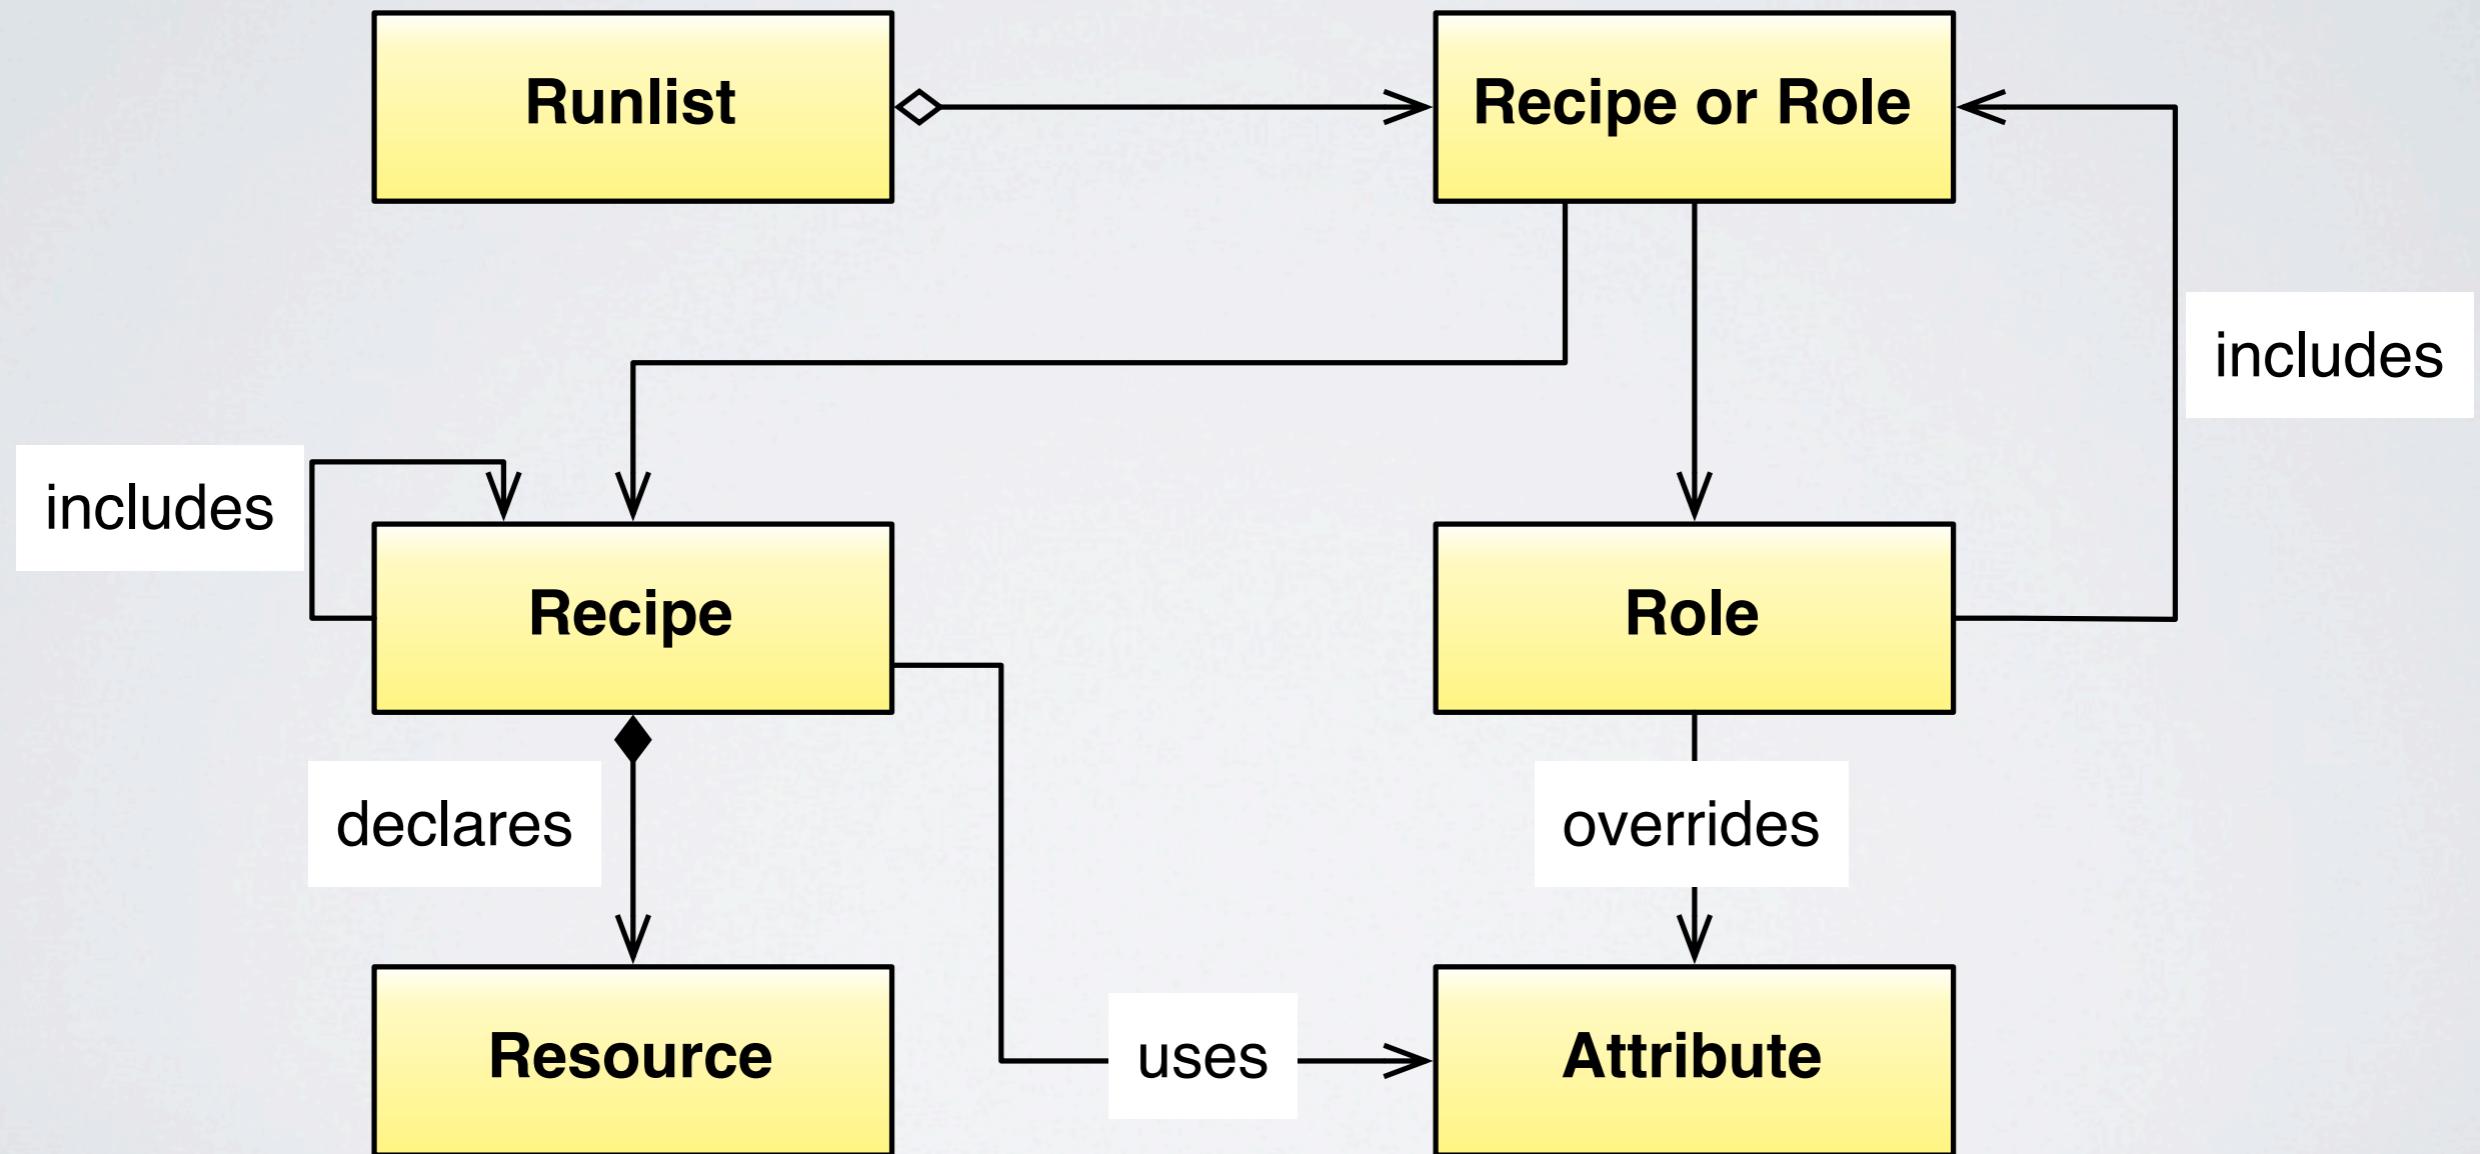
```
include Opscode::Mysql::Database

action :create_db do
  unless exists?
    begin
      Chef::Log.info "mysql_database: Creating database..."
      db.query("create database #{new_resource.database}")
      new_resource.updated_by_last_action(true)
    ensure
      db.close
    end
  end
end
```

How do cookbooks get  
applied to servers



# Runlist Executes On Node



# Example Runlist (from a role file)

```
name "dev_env_bamboo"
description "Continuous Integration server"

run_list "recipe[aws]", "recipe[xfs]",
"recipe[bamboo]",
"recipe[bamboo::crowd_authentication]",
"recipe[bamboo::ebs_volume]",
"recipe[build_dependencies]",
"recipe[maven]", "recipe[maven3]",
"recipe[bamboo::post_maven.bamboo]",
"recipe[bamboo::start]",
"recipe[iptables::ci_instance]"
```

Bamboo has a one time, setup wizard  
in the admin GUI.

How do you automate that



Recreate the results, not the process.

Recreate the results, not the process.

# Recreate the results, not the process.

- I. Install once by hand, stop before setup wizard.

# Recreate the results, not the process.

1. Install once by hand, stop before setup wizard.
2. Make a tarball.

# Recreate the results, not the process.

1. Install once by hand, stop before setup wizard.
2. Make a tarball.
3. Go through setup wizard.

# Recreate the results, not the process.

1. Install once by hand, stop before setup wizard.
2. Make a tarball.
3. Go through setup wizard.
4. Compare resulting files with snapshot.

# Recreate the results, not the process.

1. Install once by hand, stop before setup wizard.
2. Make a tarball.
3. Go through setup wizard.
4. Compare resulting files with snapshot.
5. Build a template to replicate the outcome.

Where do you keep installation  
binaries



Control, integrity, efficiency

# Control, integrity, efficiency

# Control, integrity, efficiency

I. Your VCS may hate binaries.

# Control, integrity, efficiency

1. Your VCS may hate binaries.
2. S3 worked for us. Private maven repo also works well.

# Control, integrity, efficiency

1. Your VCS may hate binaries.
2. S3 worked for us. Private maven repo also works well.
3. Vendor download might also work, but beware slipstreaming.

# Remote file resource

```
remote_file "crowd" do
  path "/tmp/crowd.tar.gz"
  source "http://downloads.atlassian.com/...
atlassian-crowd-2.3.3.tar.gz"
  checksum "dcc486625e96925..."
end
```

# Have a sandbox!

Where do I find cookbooks



Many available on github

[cookbooks.opscode.com](http://cookbooks.opscode.com)

# Download at the Shell

```
$ knife cookbook site download postgresql
Downloading postgresql from the cookbooks site ...
Cookbook saved: /Users/mtnygard/cookbooks/postgresql...
$ tar xvzf postgresql-0.11.1.tar.gz
x postgresql/
x postgresql/attributes/
x postgresql/metadata.json
x postgresql/metadata.rb
x postgresql/README.rdoc
x postgresql/recipes/
x postgresql/templates/
x postgresql/templates/default/
x postgresql/templates/default/debian.pg_hba.conf.erb
...
...
```

# One other cool knife trick

```
$ gem install knife-ec2
Fetching: excon-0.6.6.gem (100%)
...
Successfully installed knife-ec2-0.5.10
6 gems installed
$ knife ec2 server create 'role[webserver]' \
  -I ami-7000f019 -f m1.small
Instance ID: i-9e7ef1fe
Flavor: m1.small
Image: ami-7000f019
Region: us-east-1
Availability Zone: us-east-1c
Security Groups: default
SSH Key: relevance_aws

Waiting for server...
```

# Throughout Development Cycle

## Local Development

# Throughout Development Cycle

Local Development

Single App (solo)

# Throughout Development Cycle

Local Development

Single App (solo)

Server Based Rollout

# Local Development



# Vagrant

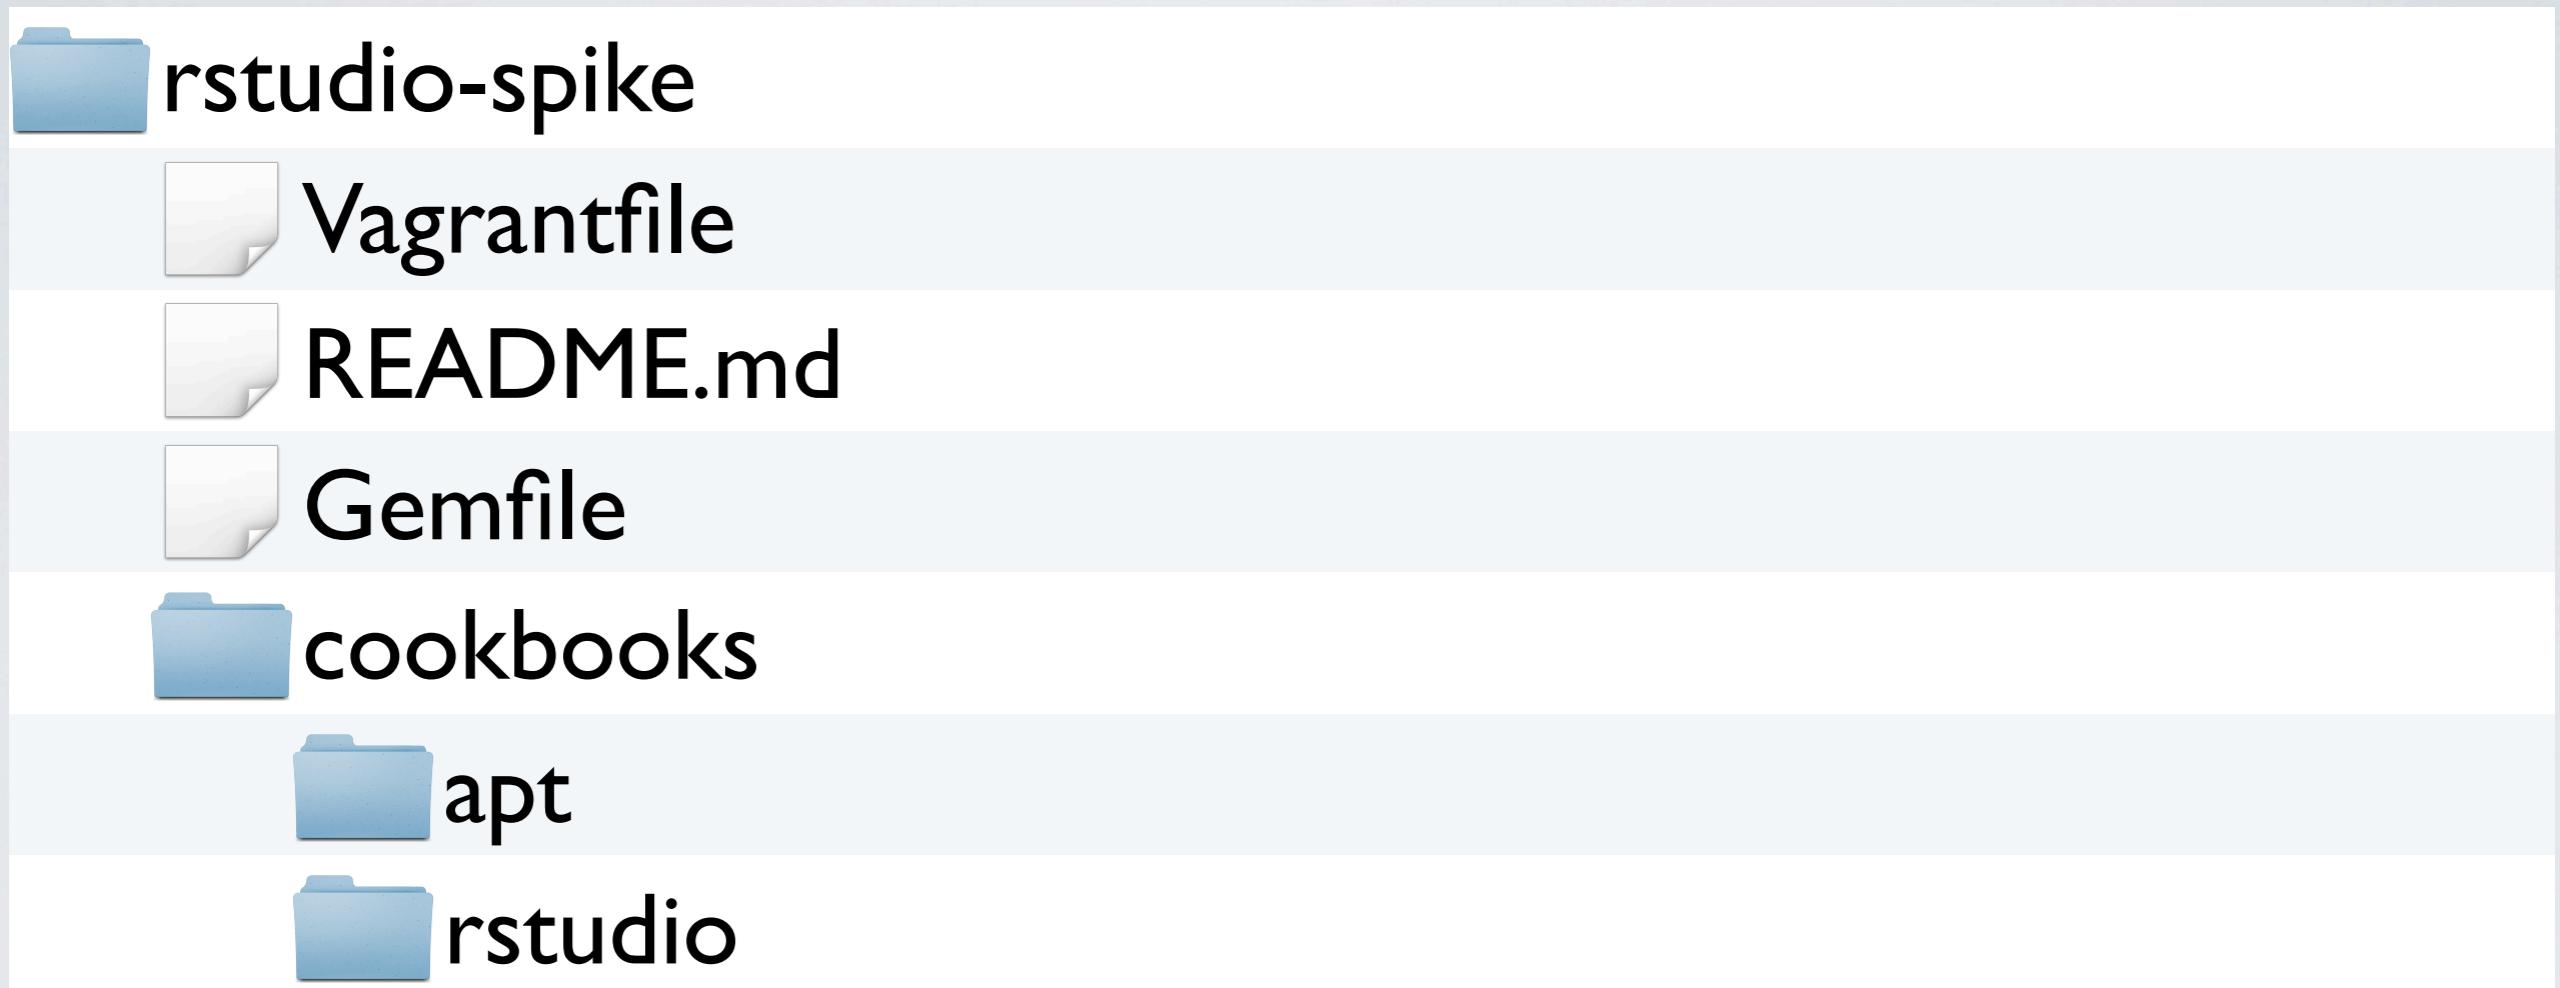


# Vagrant



Builds virtual machines from a text file and some Chef cookbooks.

# Vagrant Project





# Vagrantfile

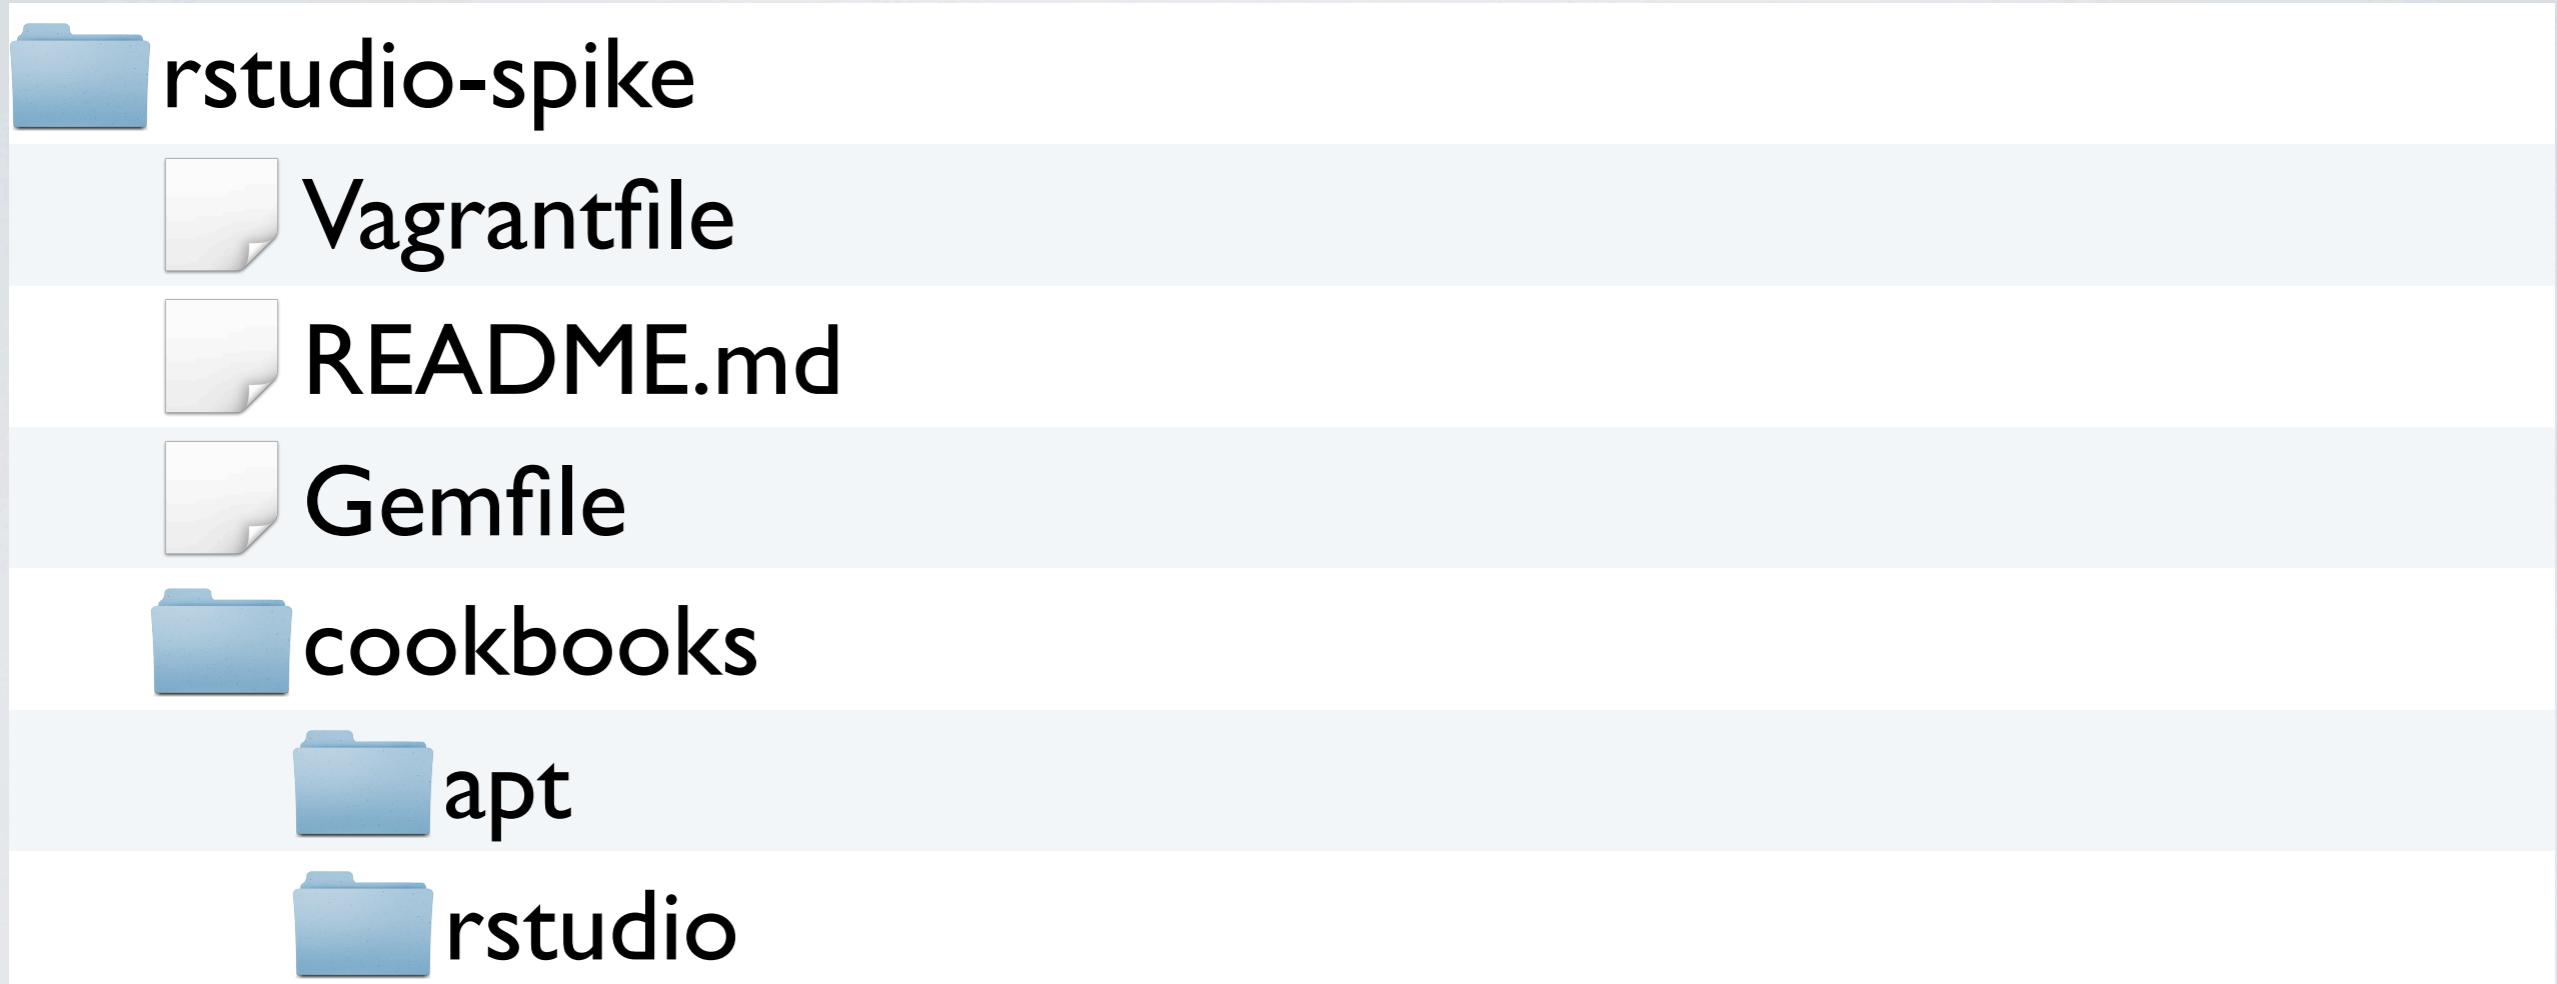
```
Vagrant::Config.run do |config|
  config.vm.box = "lucid64"
  config.vm.box_url =
    "http://files.vagrantup.com/lucid64.box"

  config.vm.forward_port "http", 8787, 8787

  config.vm.provision :chef_solo do |chef|
    chef.cookbooks_path = "cookbooks"
    chef.add_recipe "rstudio"
    chef.add_recipe "mysql::server"

    chef.json.merge!({ :mysql =>
      { :server_root_password => "foo" } })
  end
end
```

# Vagrant Project



# Benefits

# Benefits

Version control dev environment

# Benefits

Version control dev environment

Share configs across team

# Benefits

Version control dev environment

Share configs across team

Recreate project or client setup

# Benefits

Version control dev environment

Share configs across team

Recreate project or client setup

Keep host environment clean

# “Base Box”

“Base Box”

OS Image

+ Vagrant configs

# Download Base Boxes

## Vagrantbox.es

Vagrant is an amazing tool for managing virtual machines via a simple to use command line interface. With a simple `vagrant up` you can be working in a clean environment based on a standard template.

These standard templates are called **base boxes**, and this website is simply a list of boxes people have been nice enough to make publicly available.

### Suggest a Box

Do you know of another high quality base box? [Let us know about it](#) and, once we've checked it out, we'll add it to the list below.

### Available Boxes

Search

[ubuntu 11.04 server i386](#) <http://dl.dropbox.com/u/7490647/talifun-ubuntu-11.04-server-i386.box>

[ubuntu 11.04 server](#) <http://dl.dropbox.com/u/7490647/talifun-ubuntu-11.04-server-amd64.box>

# Make Base Boxes

The screenshot shows a GitHub repository page for the user 'jedi4ever' with the repository name 'veewee'. The 'Source' tab is selected, showing basic repository statistics: 2 pull requests, 34 issues, and 2 wiki pages. Below this, there are dropdown menus for 'Switch Branches' (3) and 'Switch Tags' (32), and a link to the 'Branch List'. A summary message 'Easing the building of vagrant boxes — [Read more](#)' is displayed. At the bottom, there are links for 'HTTP' and 'Git Read-Only' access, with the URL <https://github.com/jedi4ever/veewee.git>. A 'Read-Only access' button is also present. A recent merge pull request is highlighted: 'Merge pull request #115 from dcarley/vbox\_sataportcount\_1' by 'jedi4ever' (2 days ago). The overall theme of the page is related to creating base boxes for Vagrant.

**jedi4ever / veewee**

**Source** Commits Network Pull Requests (2) Issues (34) Wiki (2) Graphs

Switch Branches (3) ▾ Switch Tags (32) ▾ Branch List

Easing the building of vagrant boxes — [Read more](#)

**HTTP** **Git Read-Only** <https://github.com/jedi4ever/veewee.git> **Read-Only** access

Merge pull request #115 from dcarley/vbox\_sataportcount\_1

**jedi4ever** authored 2 days ago

**veewee /**

# Some Recommendations

Know your image;  
build your own base boxes.

Avoid multiple Chef and Vagrant copies.

Use RVM with .rvmrc in your Vagrant directory.

Reuse cookbooks between vagrant  
boxes and real deployments.

Share the project directory with the  
VM.

# Solo Deployment



# Chef Solo

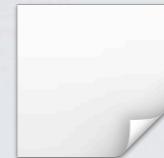
# Chef Solo

Runs cookbooks,  
with attributes,  
on a node.

# chef-solo

```
$ sudo chef-solo -c solo.rb -j node.json
...
... INFO: *** Chef 0.10.4 ***
... INFO: Setting the run_list to ["recipe[rstudio]"] from JSON
... INFO: Run List is [recipe[rstudio]]
... INFO: Run List expands to [rstudio]
... INFO: Starting Chef Run for lucid64.hsd1.ca.comcast.net.

...
... INFO: Processing package[r-base] action install
(rstudio::default line 18)
                                -- many lines of output --
... INFO: Processing dpkg_package[rstudio-server] action install
(rstudio::default line 46)
... INFO: dpkg_package[rstudio-server] installed version 0.94.84
... INFO: Chef Run complete in 238.033898 seconds
... INFO: Running report handlers
... INFO: Report handlers complete
```



## solo.rb

```
file_cache_path "/var/chef-solo"
cookbook_path ["/var/chef/cookbooks",
               "/var/chef/site-cookbooks"]
```



## node.json

```
{  
  "mysql": {  
    "server_root_password": "foo"  
  },  
  "run_list": ["recipe[rstudio]"]  
}
```

# Benefits

# Benefits

## Self-contained bundle

# Benefits

Self-contained bundle

Can run from tarball at URL

# Benefits

Self-contained bundle

Can run from tarball at URL

Easy setup; no extra infrastructure

# Some Recommendations

Rubyists: deliver cookbooks with  
capistrano, use chef-solo to execute.

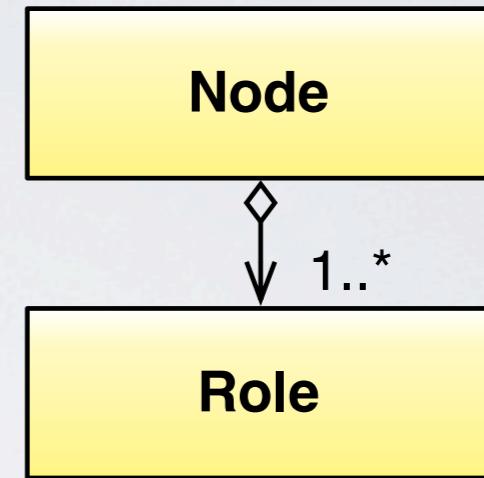
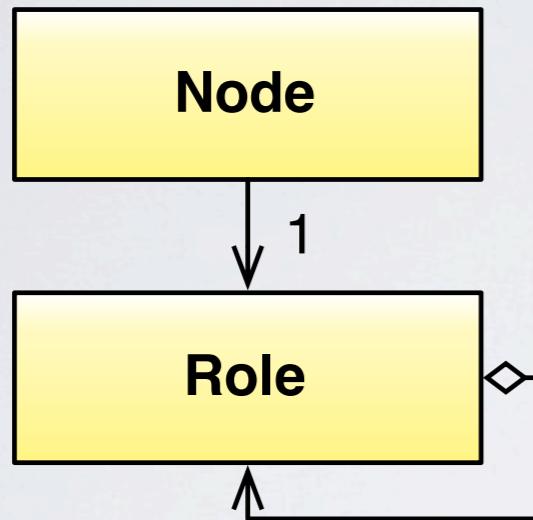
Everyone: start with chef-solo before  
you deal with chef server.

Make a shared repository for  
cookbooks.

Use attributes to factor out  
environment and tuning.

Use roles to set parameters for an environment.

# Factoring Roles



## Compose Role

Runlist for role includes other roles

Use for applications

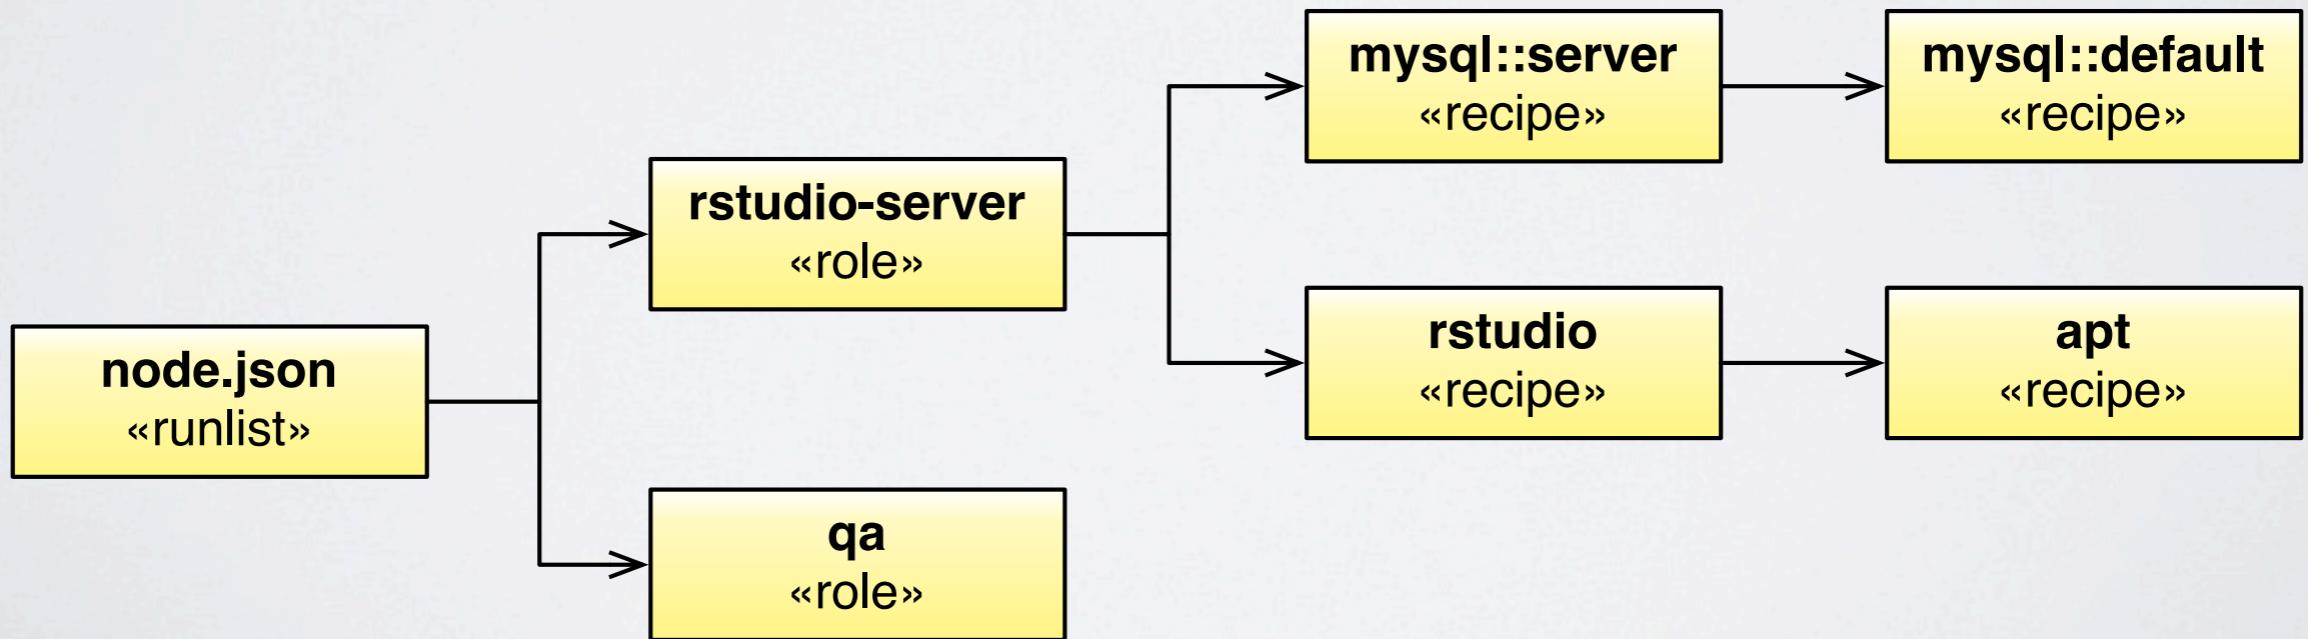
## Mix in Role

Runlist for role includes other roles

Use for environments

# Example: Factoring Roles

```
{  
  "run_list": ["role[rstudio-server]",  
               "role[qa]"]  
}
```



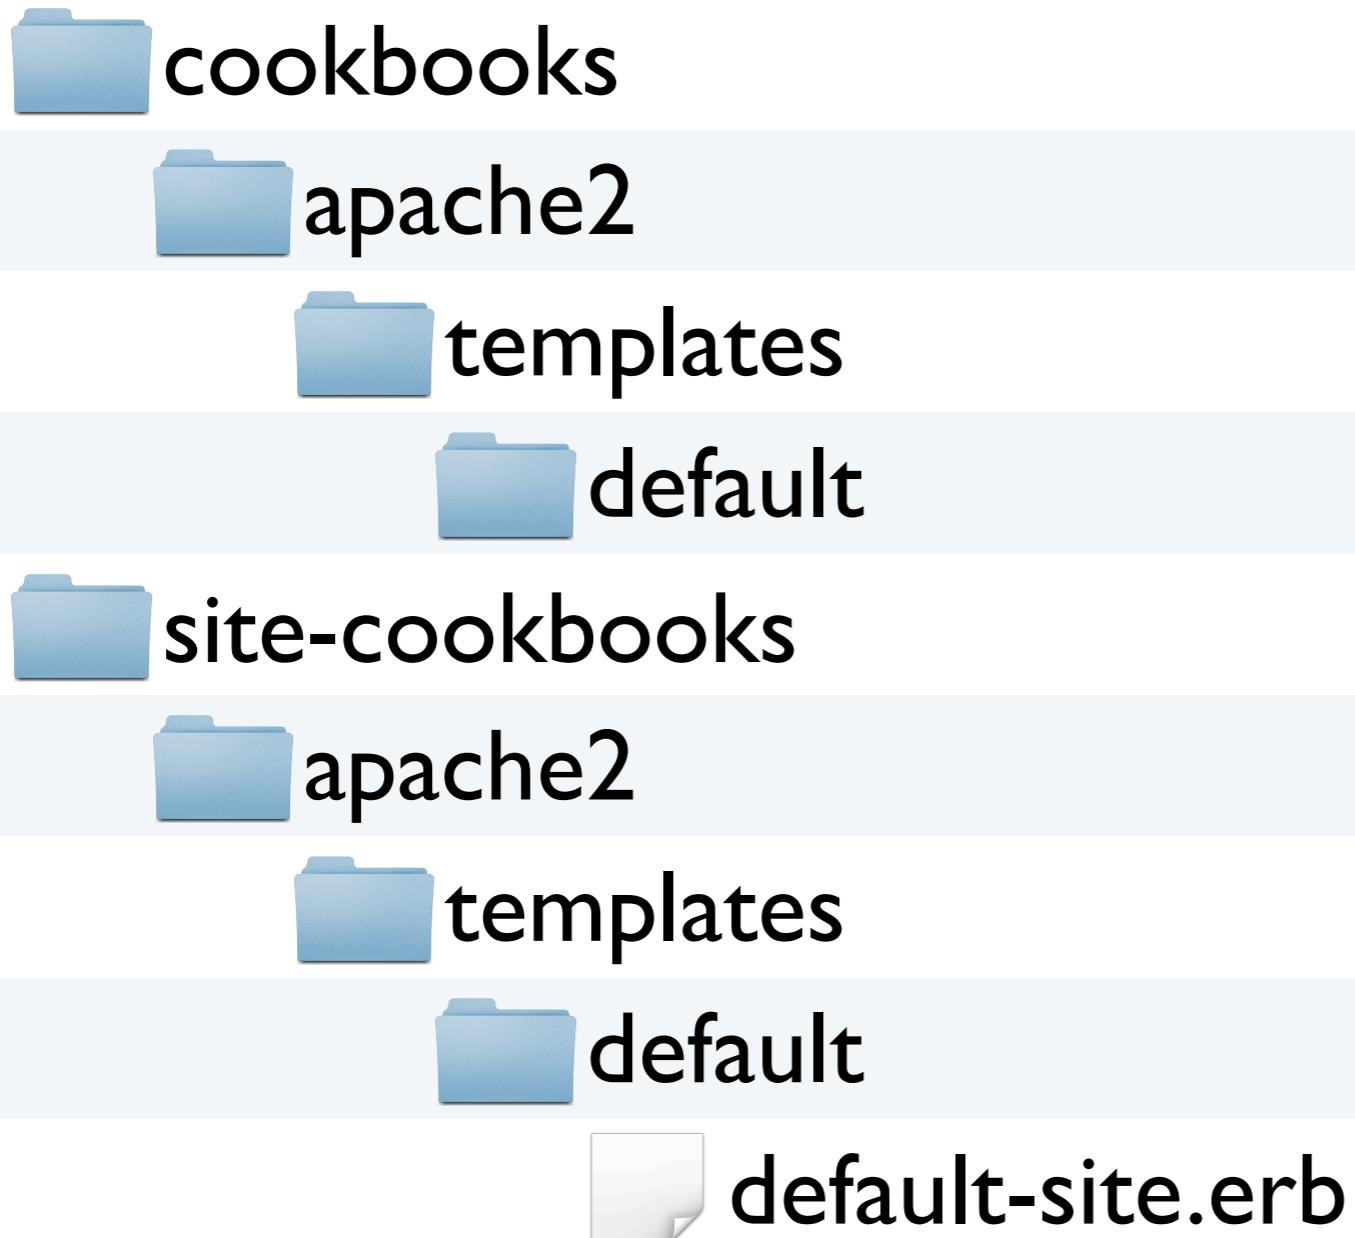
Separate off-the-shelf cookbooks  
from locally maintained cookbooks

# Example: Isolating Local Cookbooks

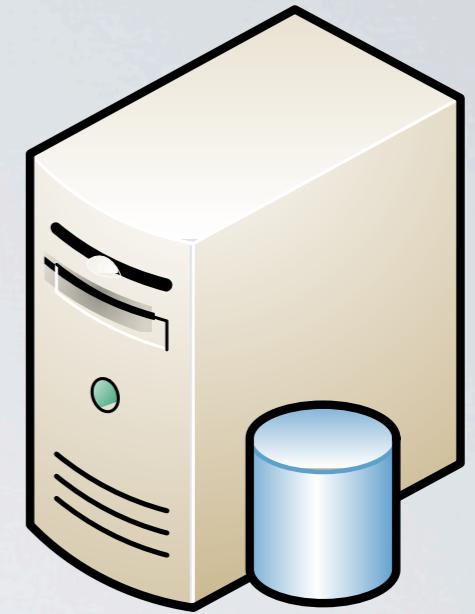
project-root

-  README.md
-  cookbooks
-  site-cookbooks
-  vagrant
-  Vagrantfile

# Example: Local Cookbooks With Overrides

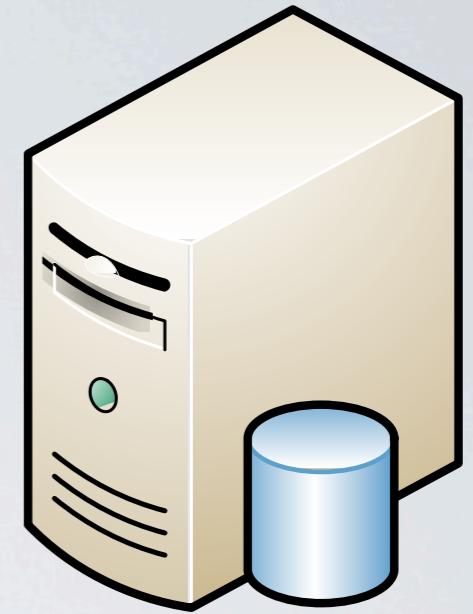


Use overrides sparingly.  
It's implementation inheritance.



# Server-Based Deployment

# Chef Server



# Chef Server

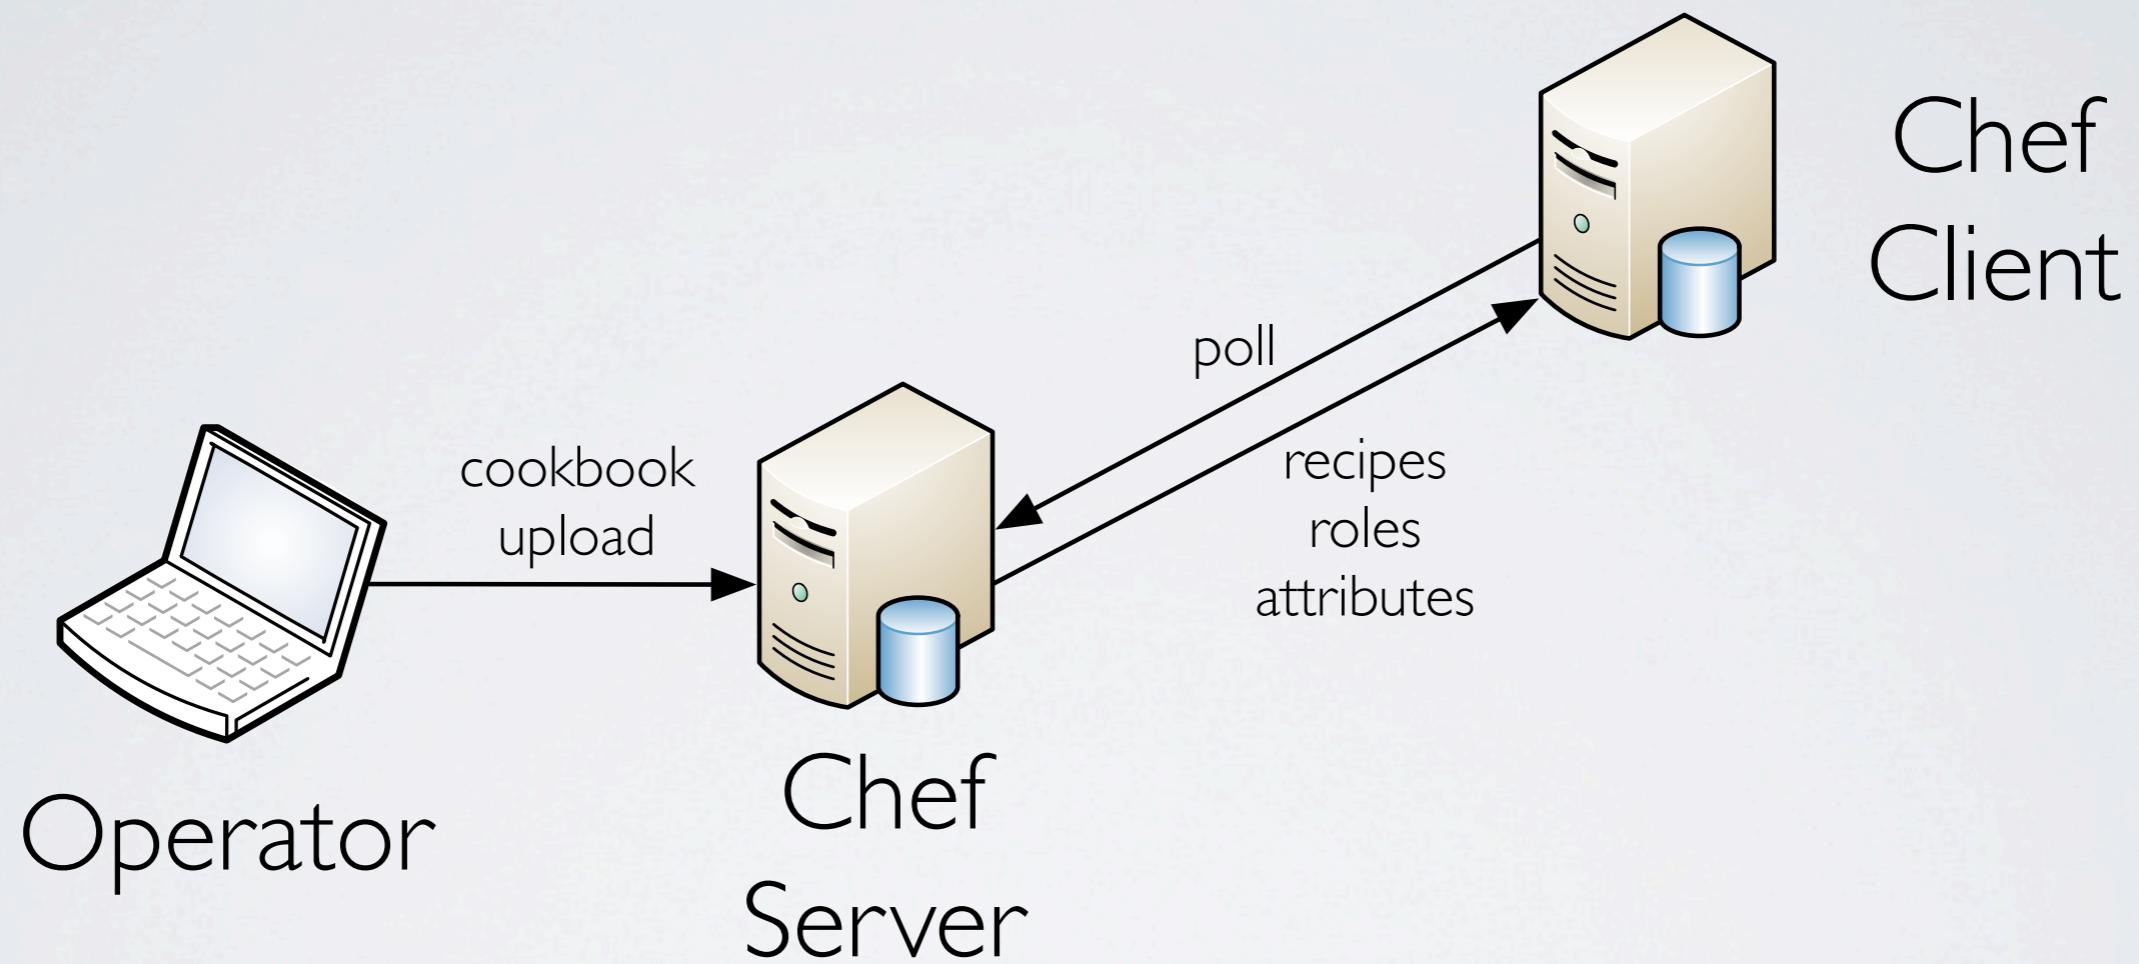
Client-server polling

Searchable via API

Authenticated and encrypted

Centralized management

# Chef Server



# Server Based Search

API access to Lucene search

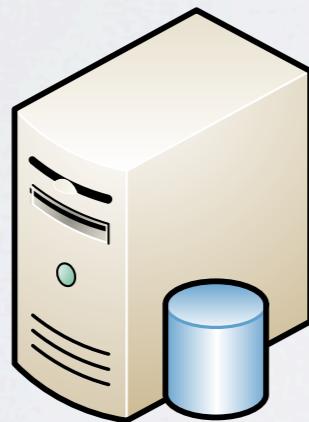
Nodes, roles, or data bags

Can be used from recipes on clients

# Example: Automatic Nagios Config

1. Search for nodes in the environment
2. Generate Nagios config files with information about nodes.

# Step 1: Chef Server Assigns Recipe “nagios::server”

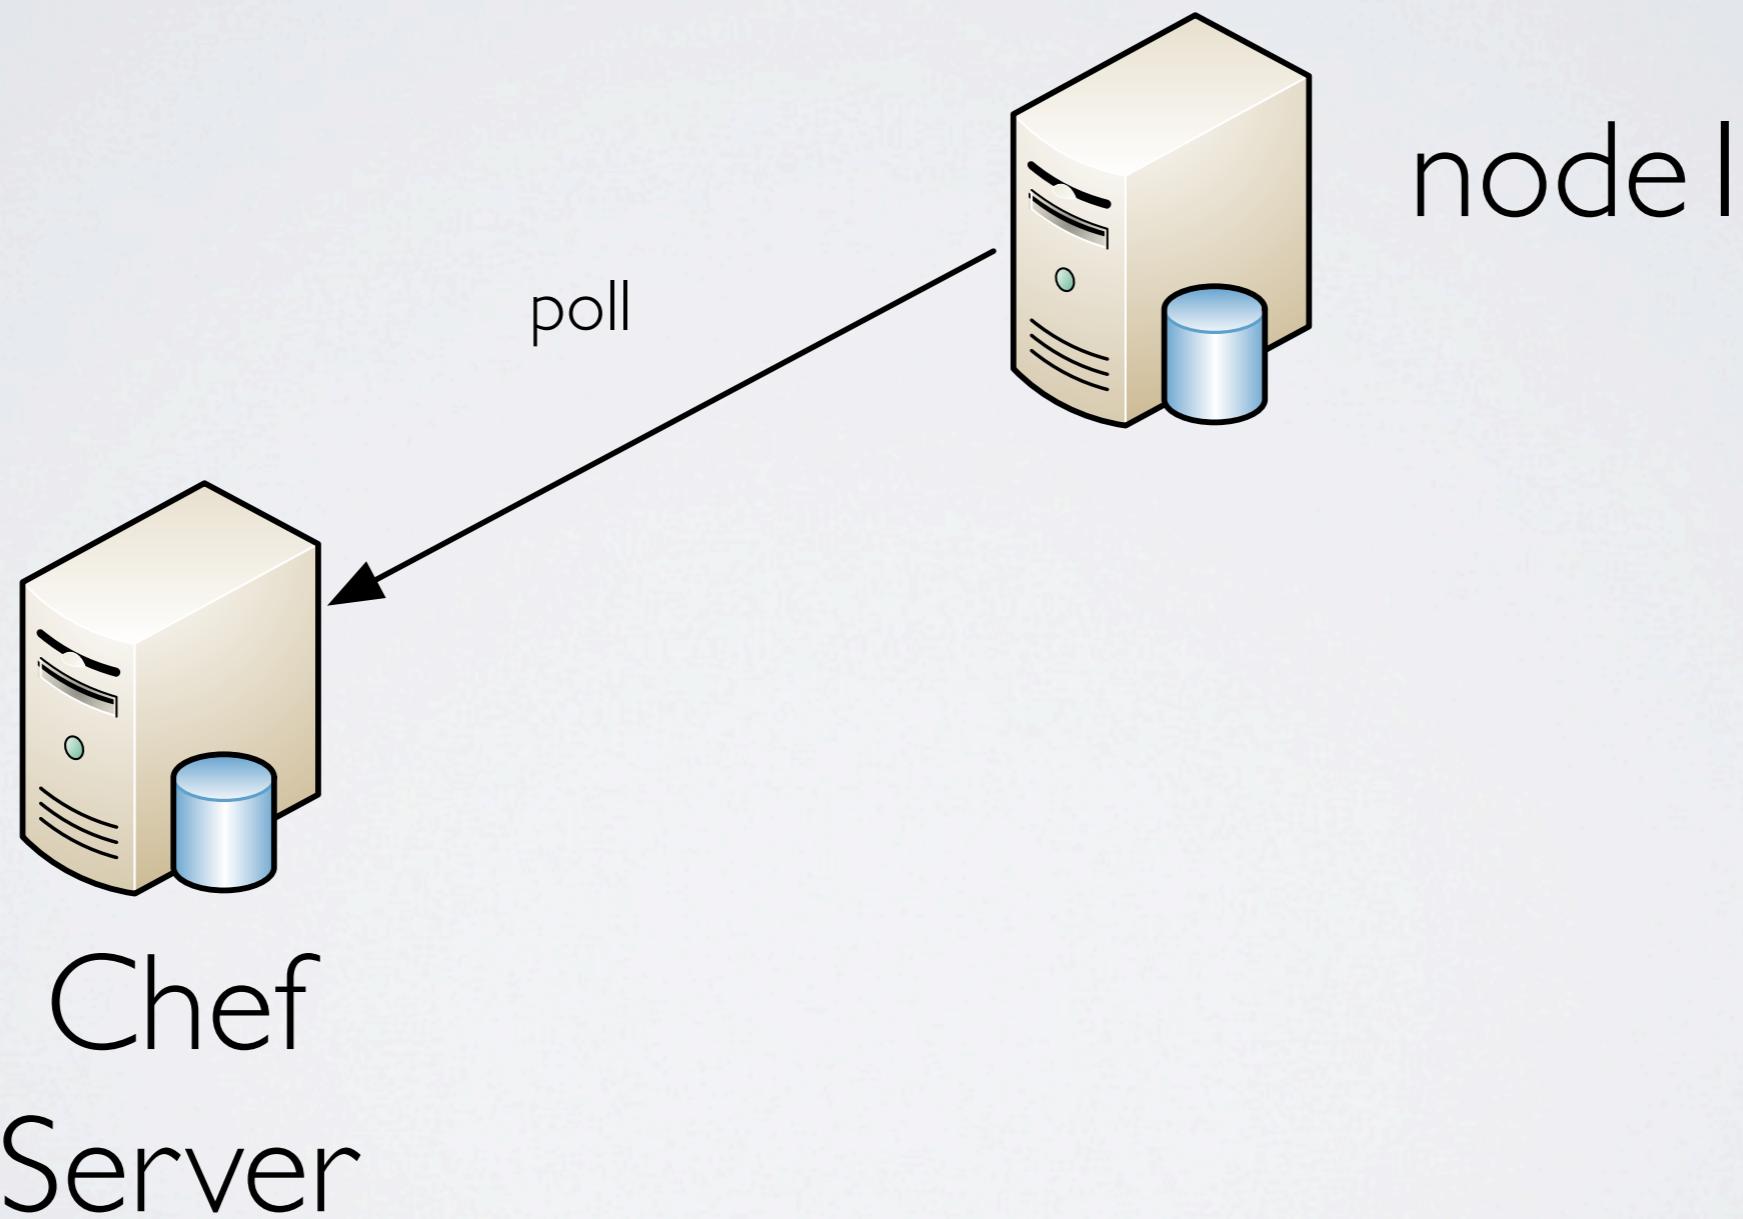


Chef  
Server

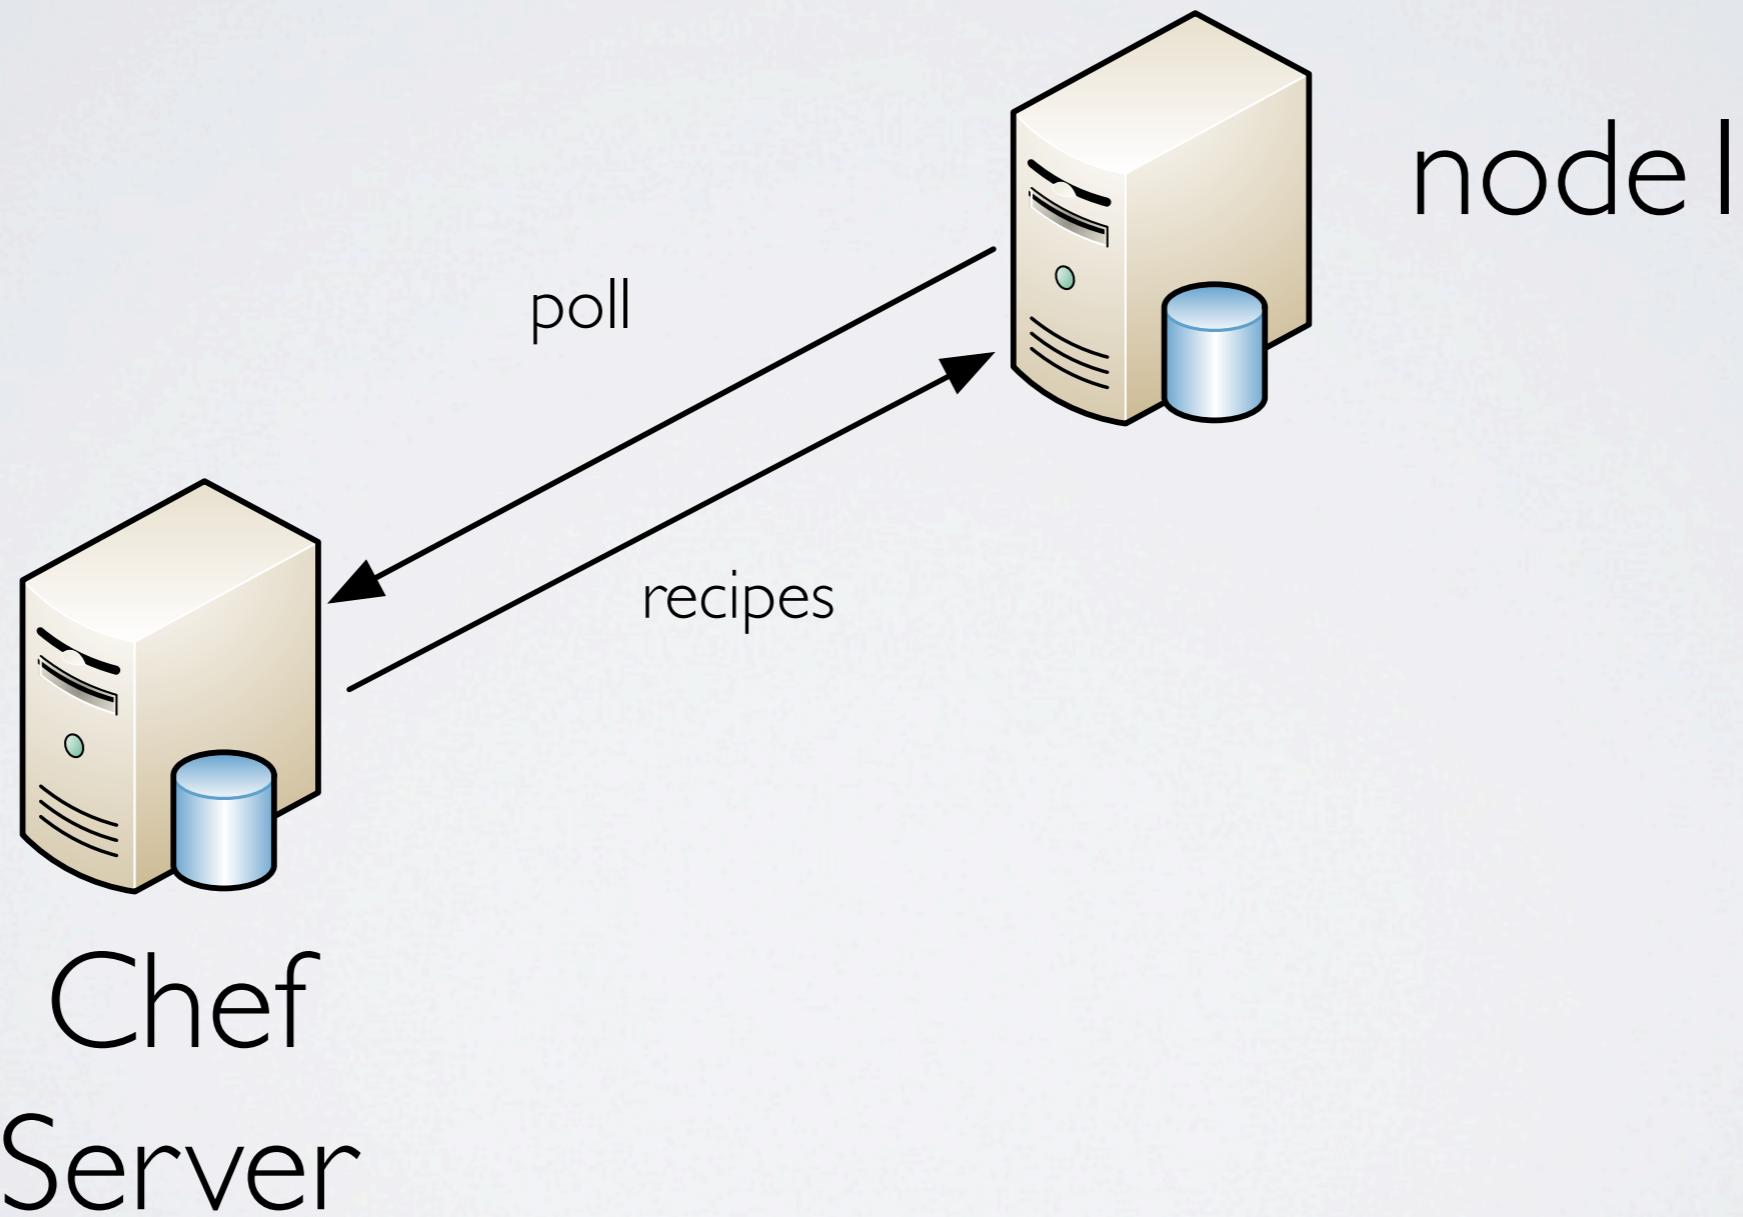


node l

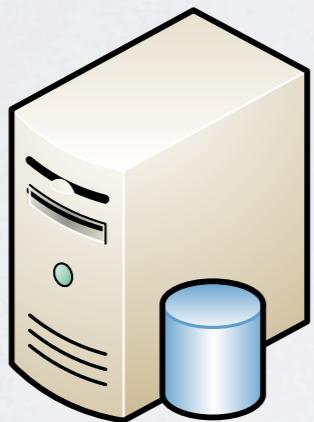
# Step 1: Chef Server Assigns Recipe “nagios::server”



# Step 1: Chef Server Assigns Recipe “nagios::server”



# Step 2: Client Downloads Cookbooks

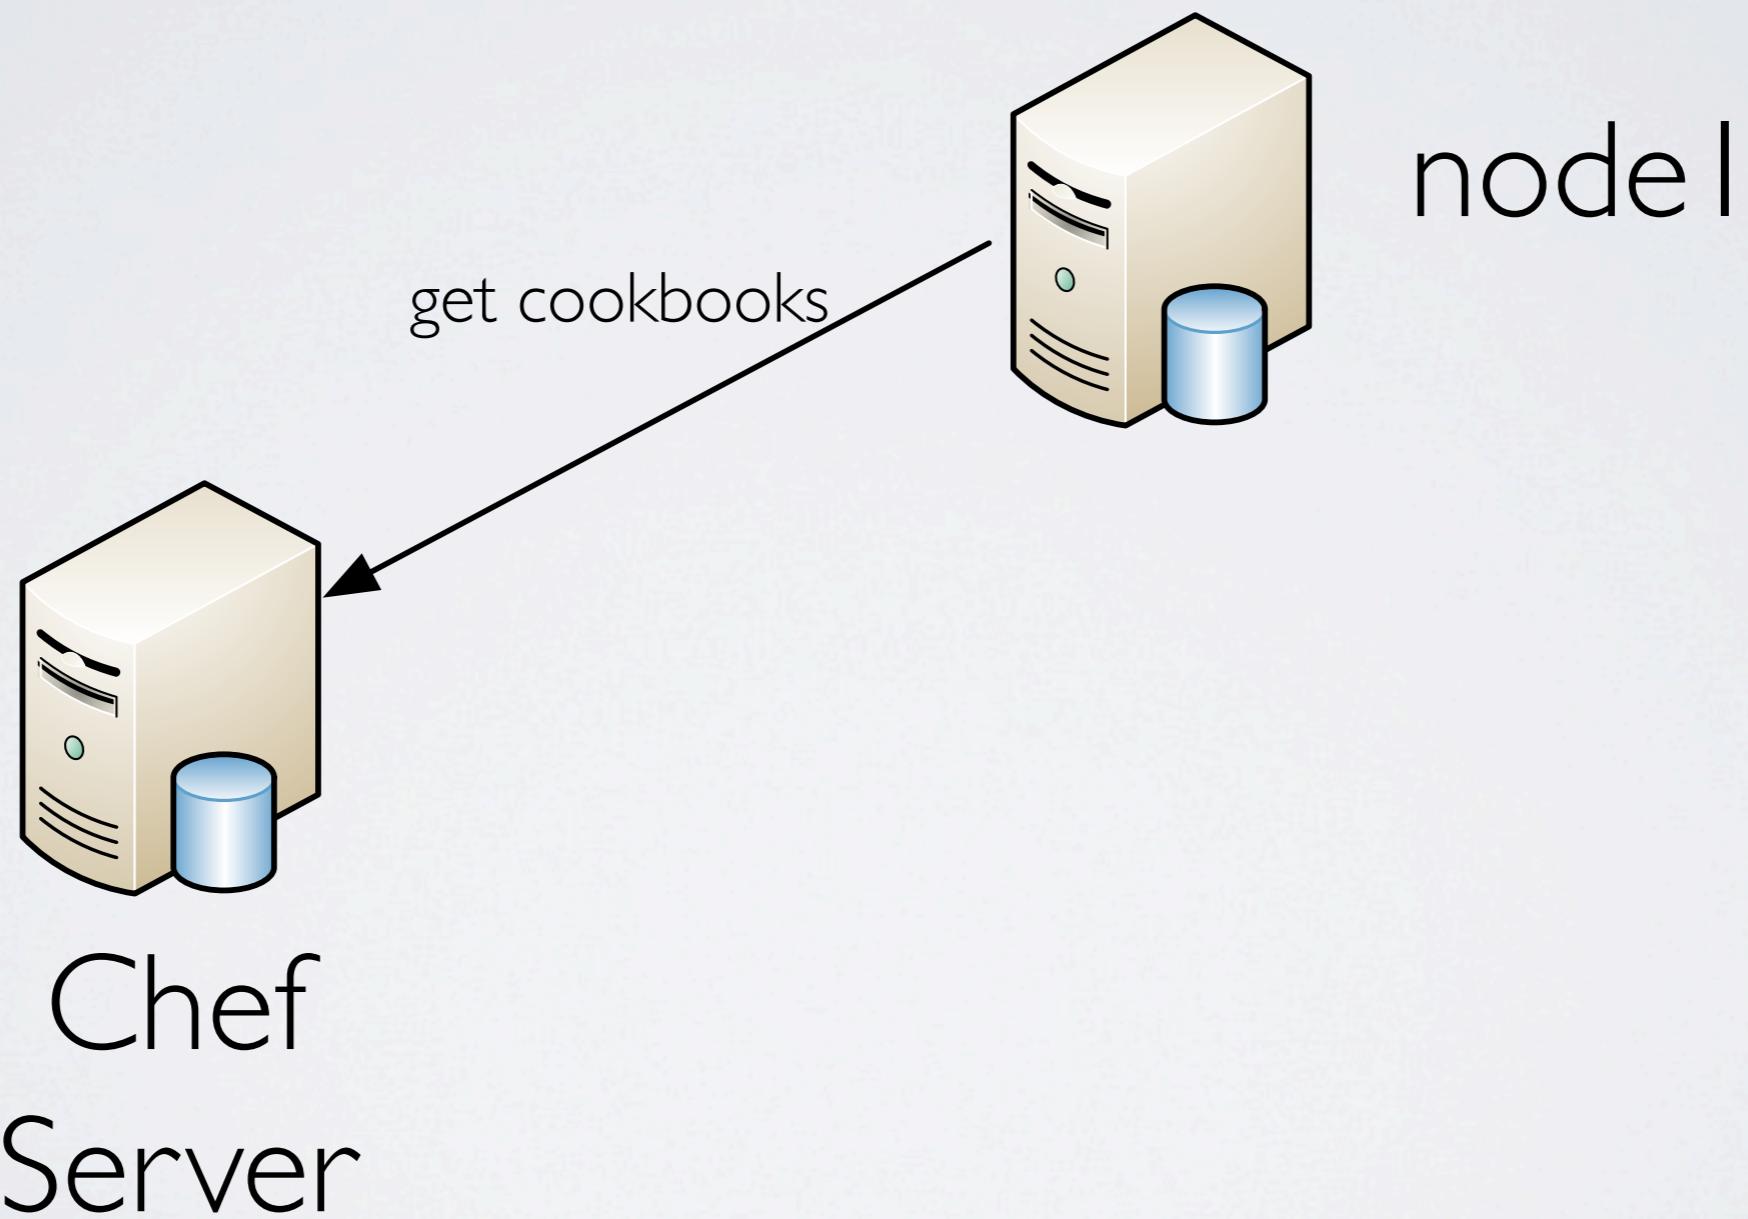


Chef  
Server

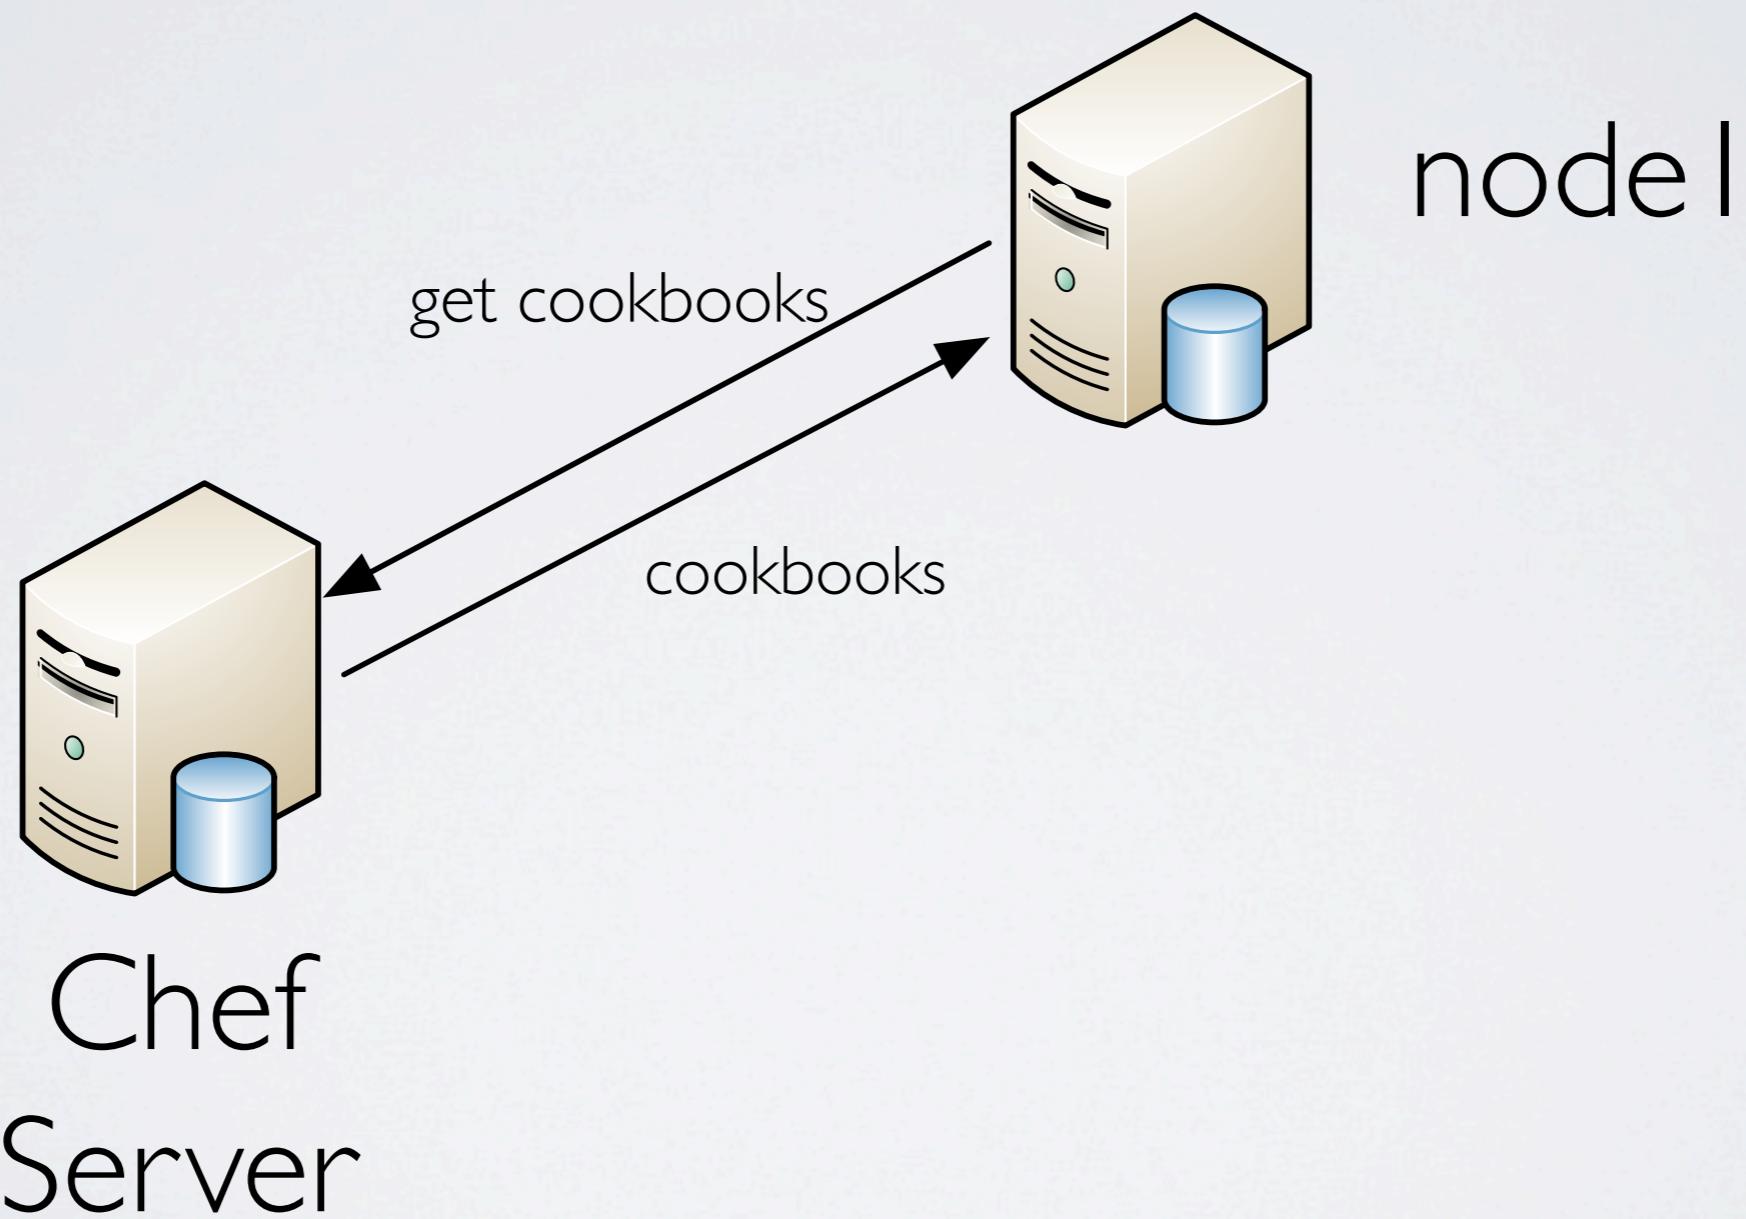


node 1

# Step 2: Client Downloads Cookbooks



# Step 2: Client Downloads Cookbooks



# Step 3: Client Executes Recipe

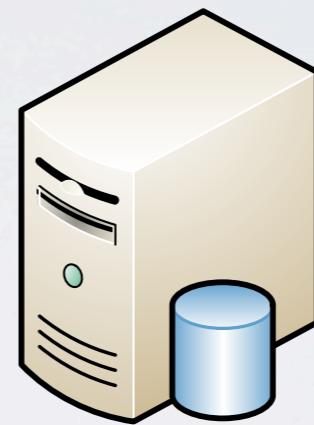
```
sysadmins = search(:users, 'groups:sysadmin')
nodes = search(:node,
  "hostname:[* TO *] AND role:#{node[:app_environment]}")

members = Array.new
sysadmins.each do |s|
  members << s['id']
end
```

# Step 3a: Client Searches Nodes

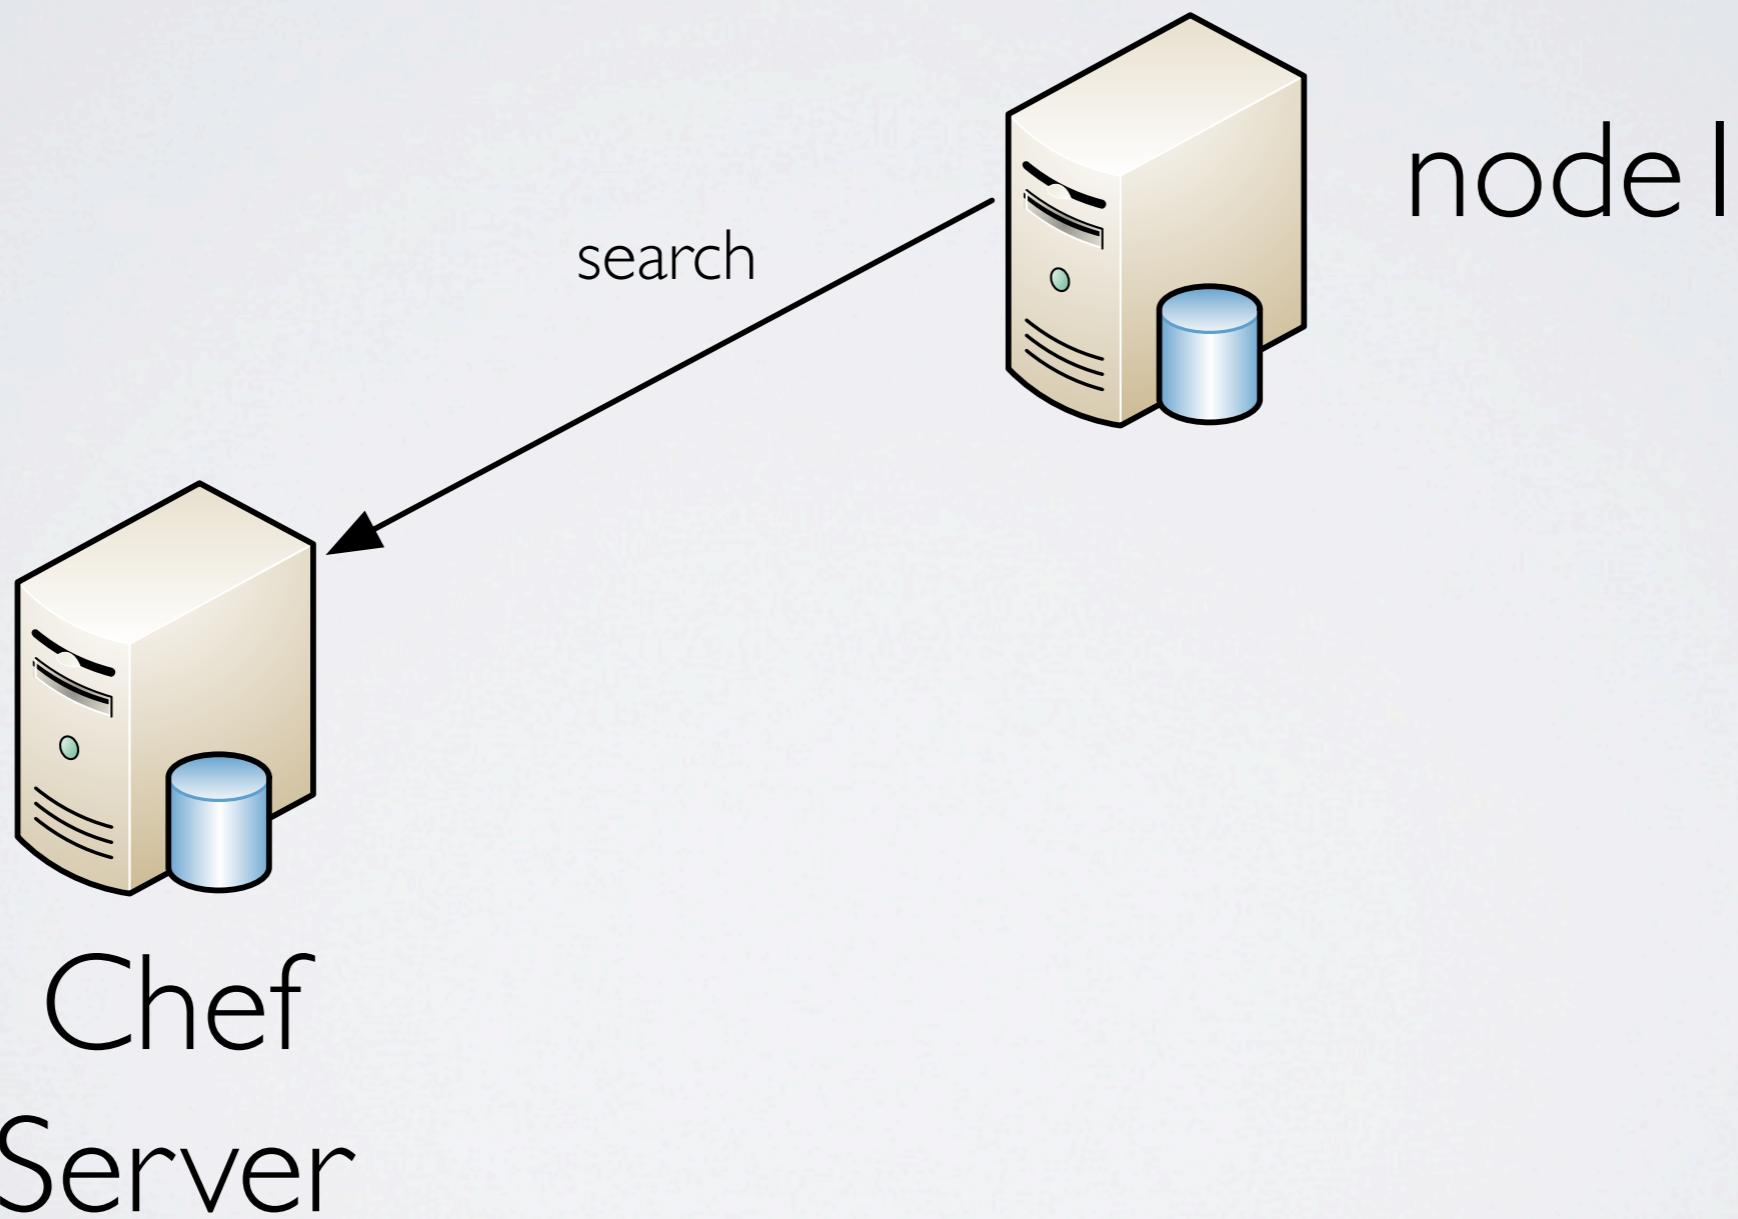


Chef  
Server

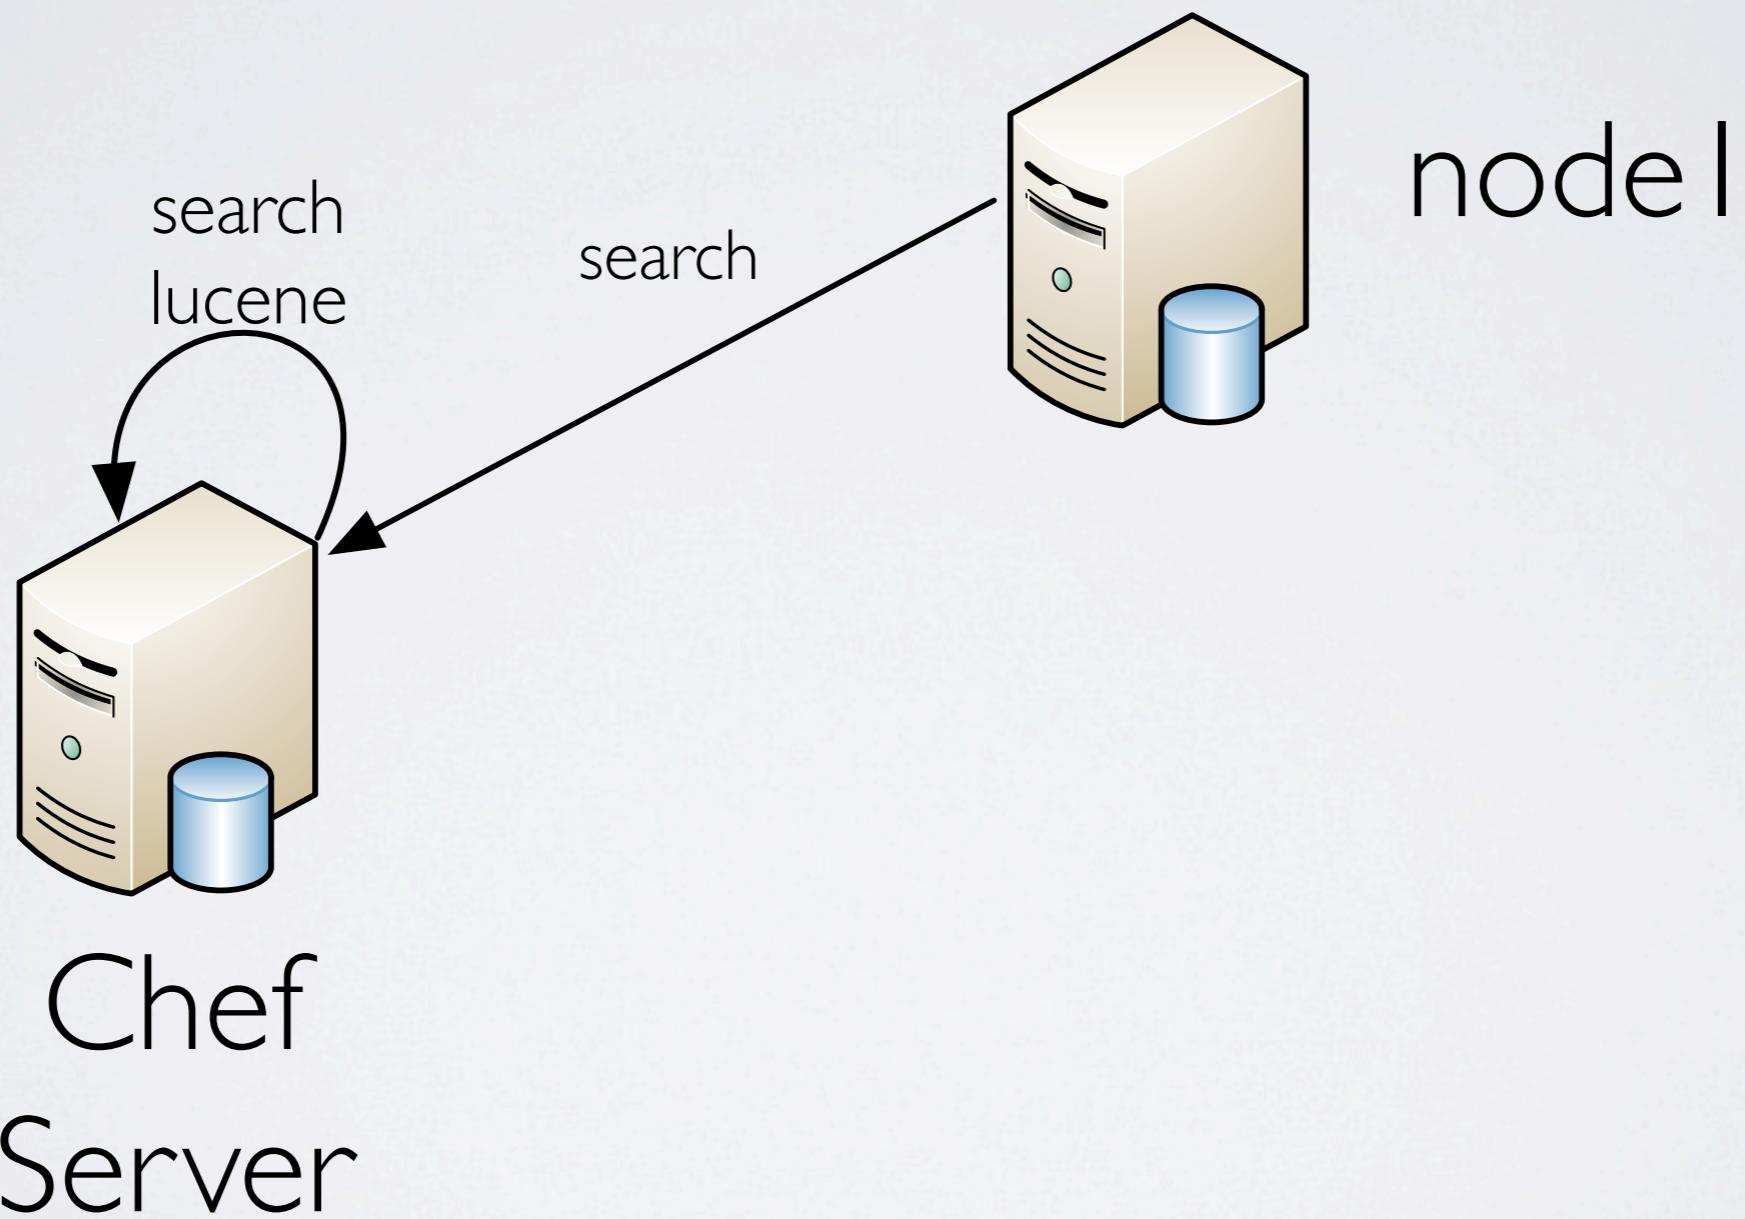


node l

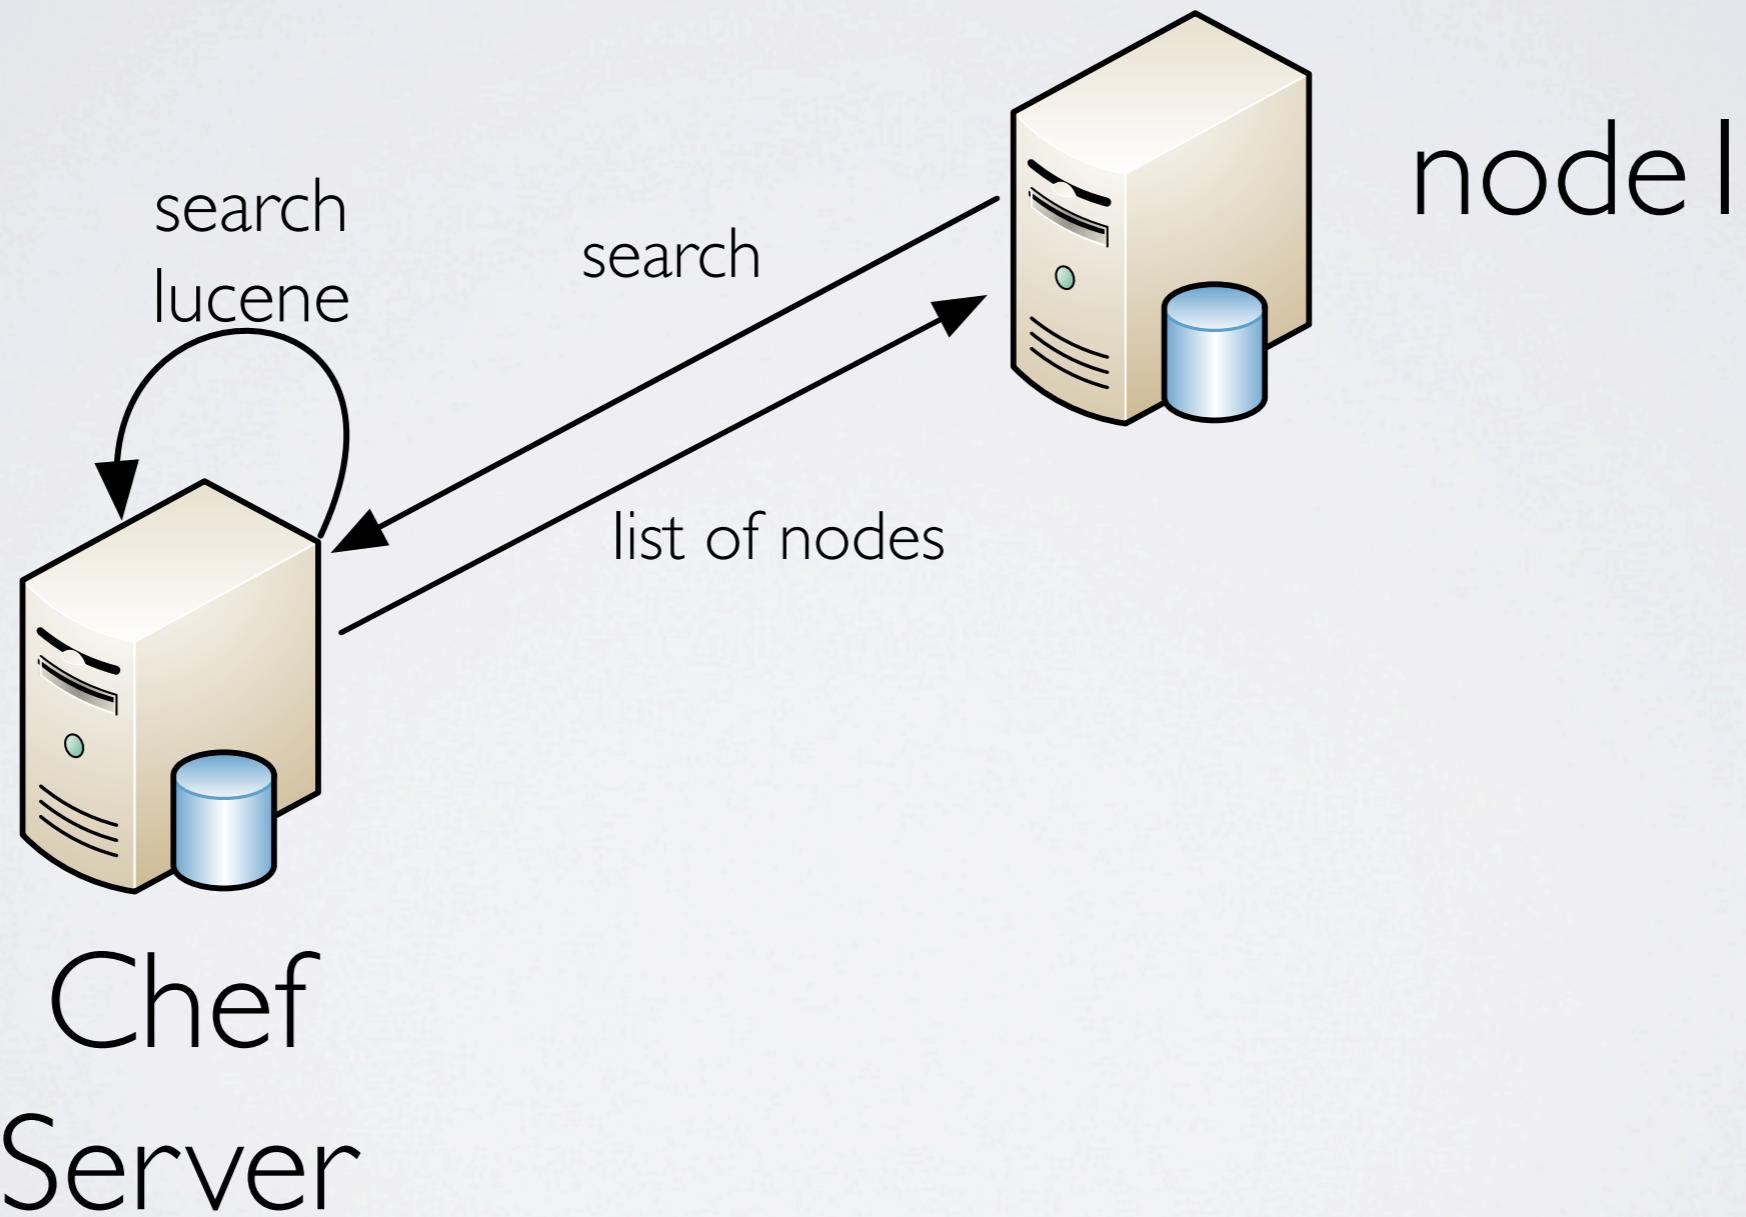
# Step 3a: Client Searches Nodes



# Step 3a: Client Searches Nodes



# Step 3a: Client Searches Nodes



# Step 3b: Client Applies Templates

```
nagios_conf "hosts" do
  variables :nodes => nodes
end
```

# Step 3b: Client Applies Templates

```
<% @nodes.each do |n| -%>
define host {
  use server
  address <%= n['ipaddress'] %>
  host_name <%= n['hostname'] %>
  hostgroups <%= n.run_list.roles.to_a.join(",") %>
}
<% end -%>
```

# Data Bags

JSON data

Held on server

Access via API

# Example: Database Credentials

```
begin
  bamboo_app = Chef::DataBagItem.load(:apps, :bamboo)
  Chef::Log.info("Loaded... apps[#{bamboo_app['id']}]")
rescue
  Chef::Log.fatal("Could not find #{:bamboo}... ")
  raise
end
```

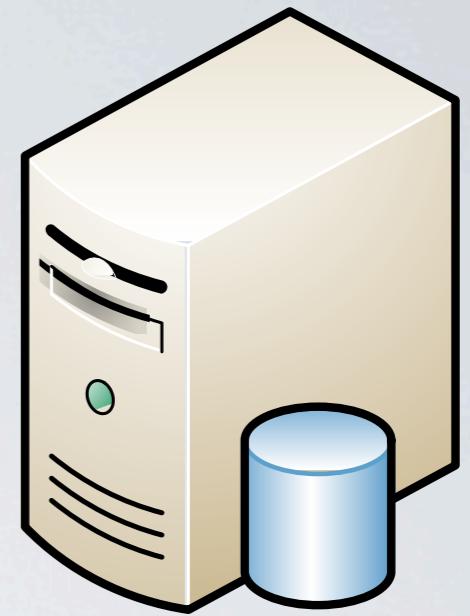
Search & data bags require Chef Server.

Breaks Vagrant and Chef Solo.

!

## Encrypted Data Bags

<https://jtimberman.posterous.com/64227128>



# Chef Server in EC2

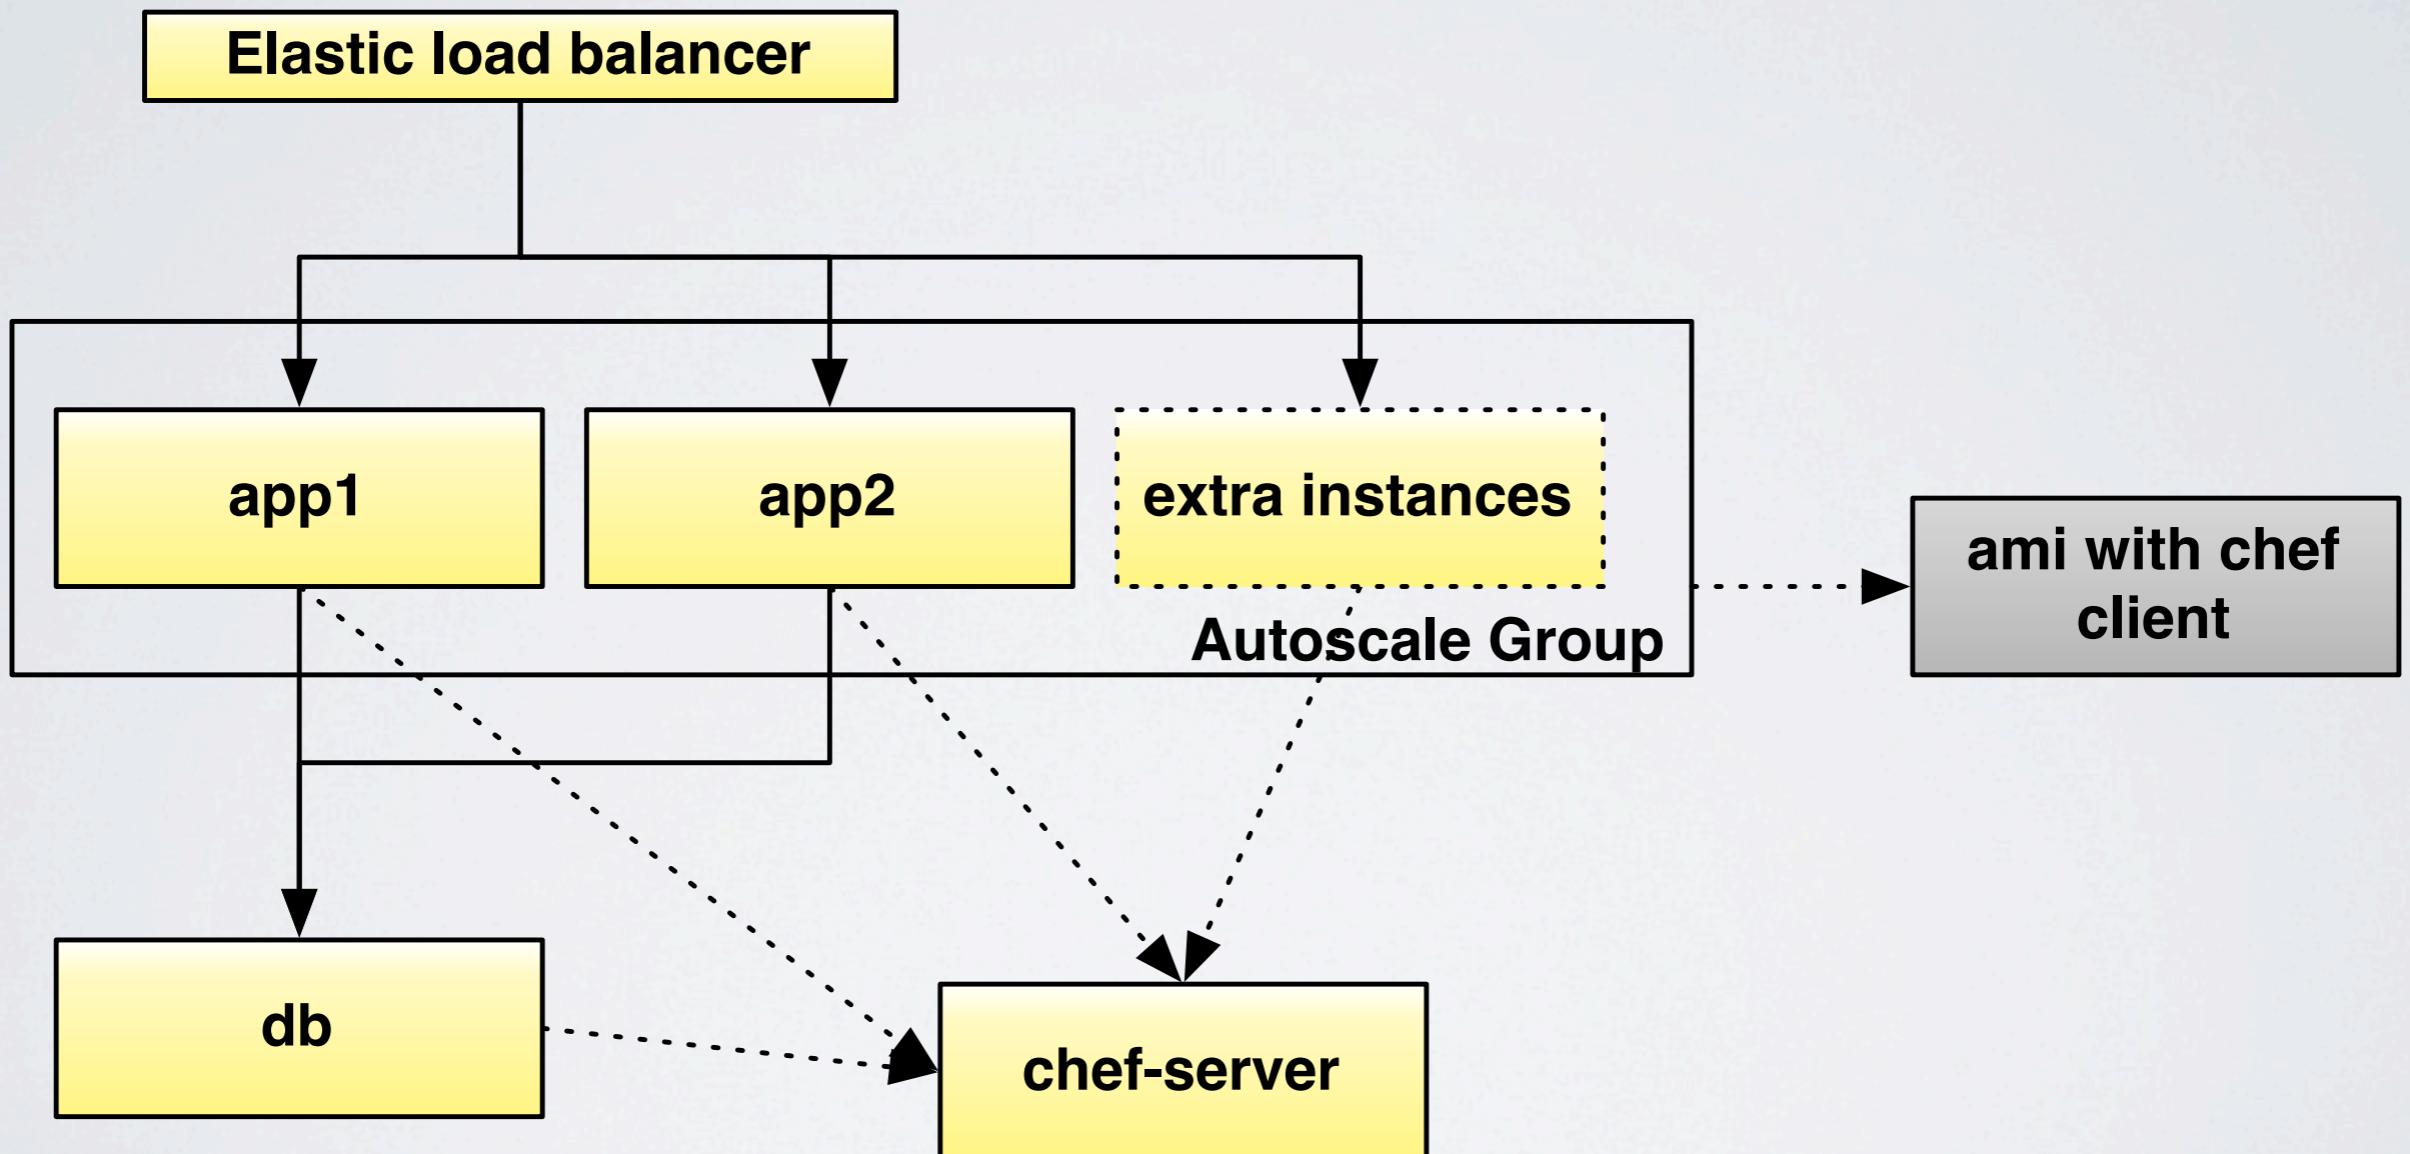
# Challenges

No long-lived identity

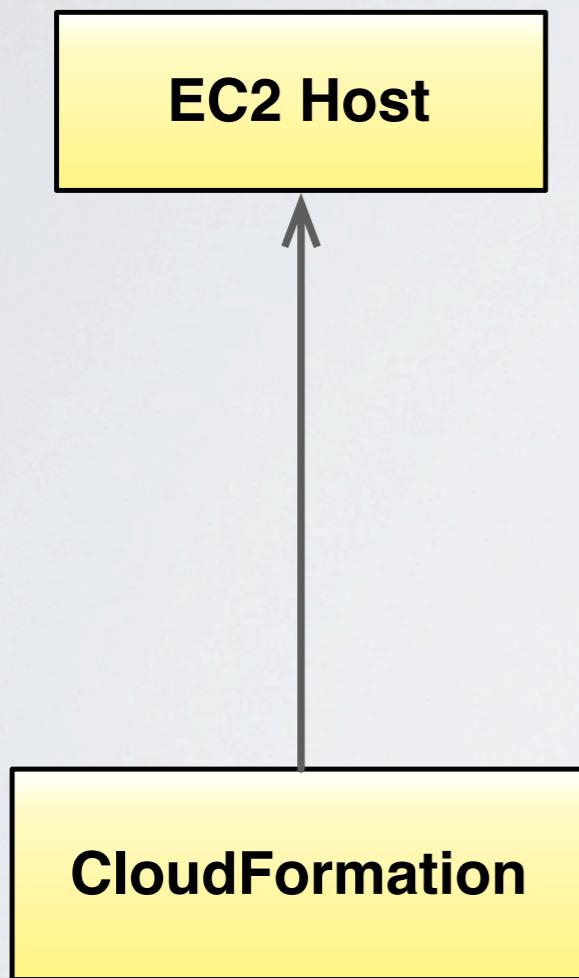
Must use EBS for server

Authentication of new nodes

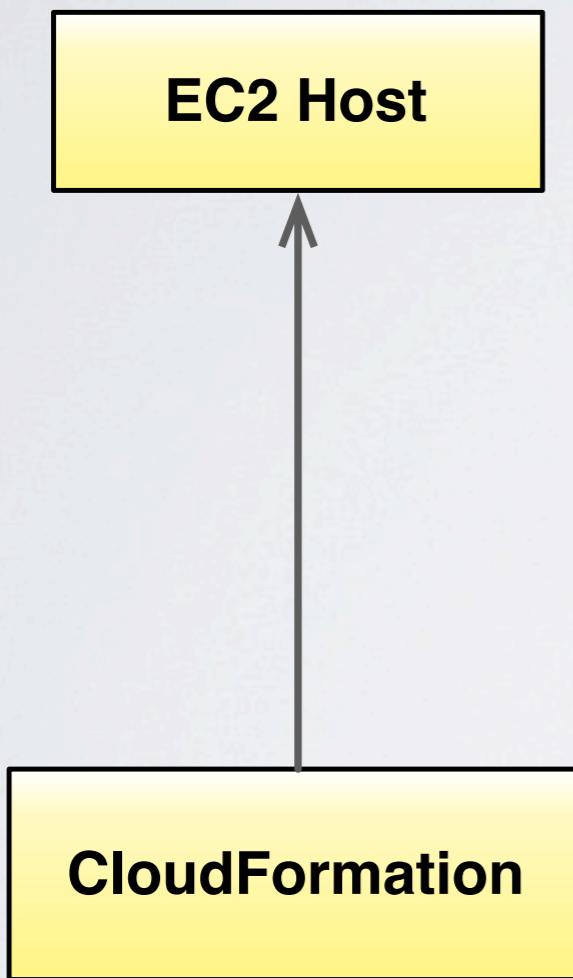
# Example: EC2 + ELB + AS



# Recent News: CloudFormation + Chef

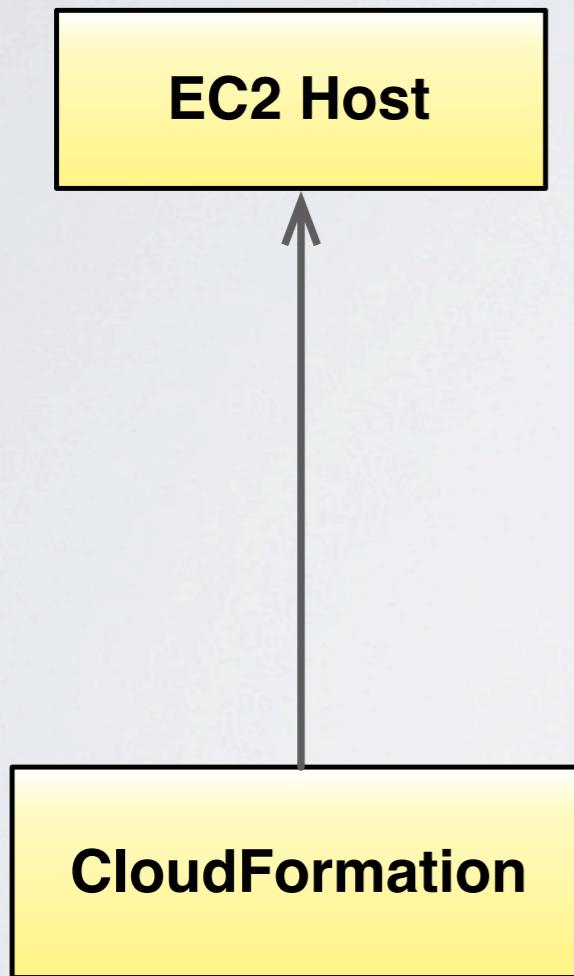


# Recent News: CloudFormation + Chef



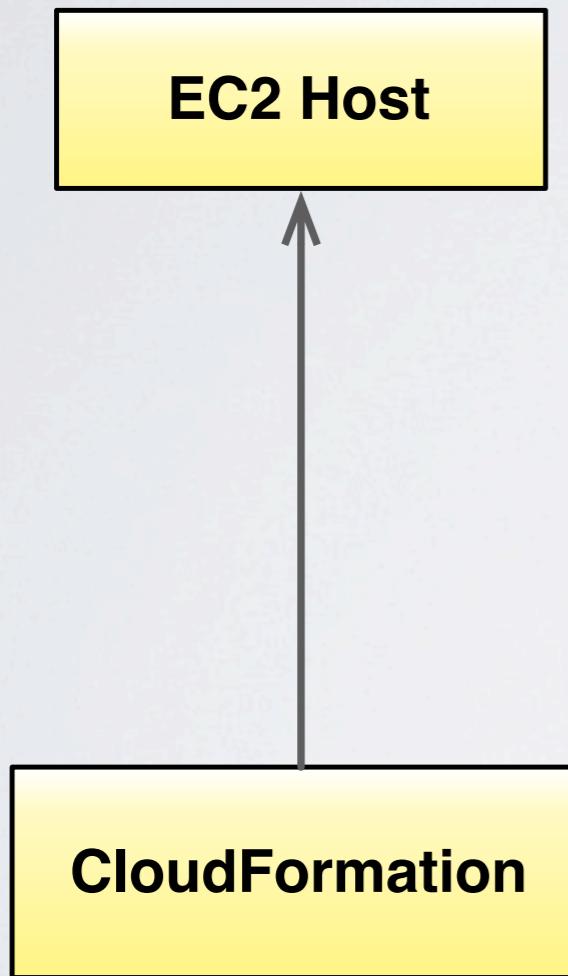
- I. create virtual machine

# Recent News: CloudFormation + Chef



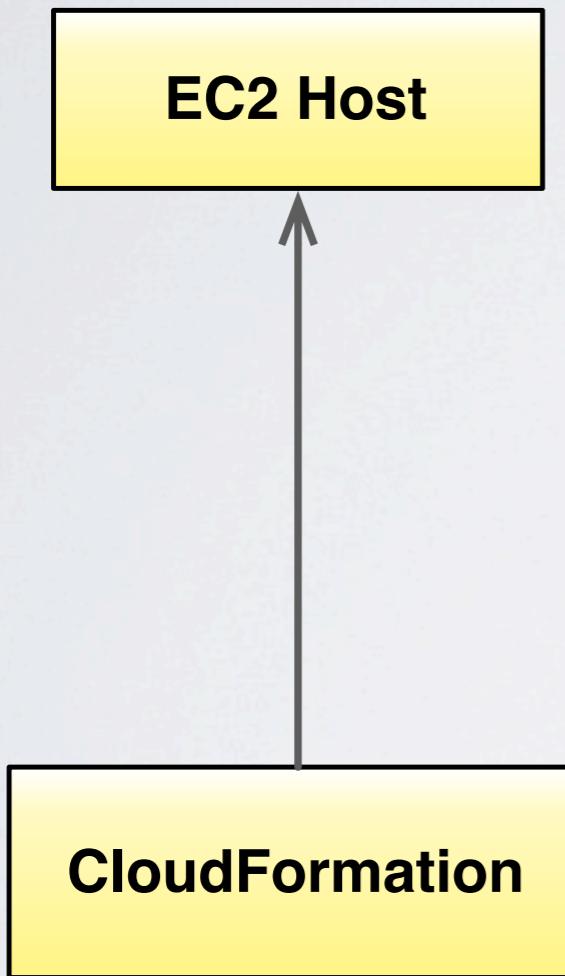
1. create virtual machine
2. yum install gcc, ruby, rubygems  
gem install chef

# Recent News: CloudFormation + Chef



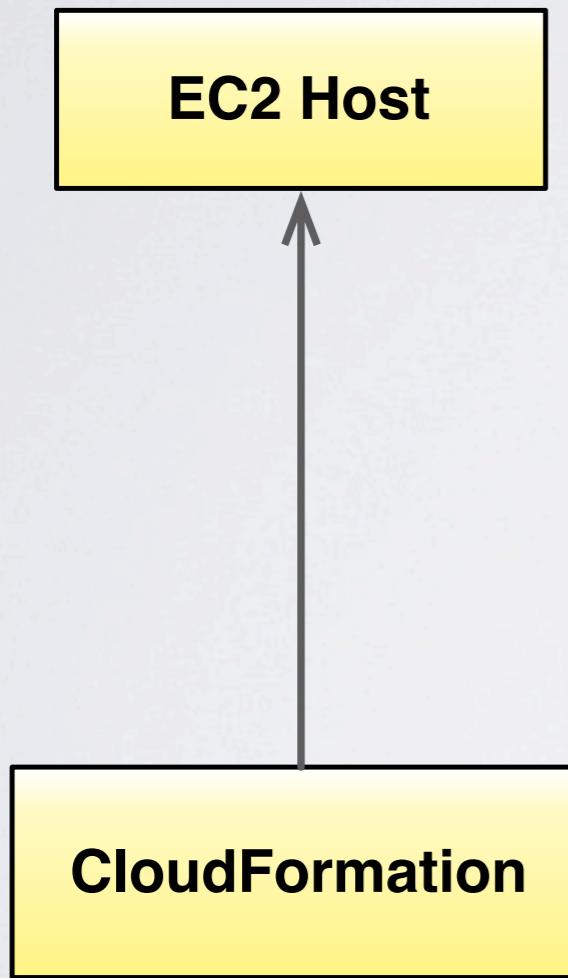
1. create virtual machine
2. yum install gcc, ruby, rubygems  
gem install chef
3. deposit solo.rb & node.json files

# Recent News: CloudFormation + Chef



1. create virtual machine
2. yum install gcc, ruby, rubygems  
gem install chef
3. deposit solo.rb & node.json files
4. run chef-solo with a URL target

# Recent News: CloudFormation + Chef



1. create virtual machine
2. yum install gcc, ruby, rubygems  
gem install chef
3. deposit solo.rb & node.json files
4. run chef-solo with a URL target
5. chef-solo fetches cookbooks &  
executes recipes

# Chef in Private Clouds

# Chef + OpenStack: Automation 2 Ways

1. Build the platform itself
2. Build nodes in the platform

# Other Options

Hosted Chef: Chef Server as a Service

[www.opscode.com](http://www.opscode.com)

# Other Options

Hosted Chef: Chef Server as a Service

Private Chef: Like Hosted Chef, as an appliance  
behind your firewall

[www.opscode.com](http://www.opscode.com)

# Resources

Chef

<http://wiki.opscode.com>

Vagrant

<http://www.vagrantup.com/>

CloudFormation

<http://bit.ly/mPFFJG>

# CHEF IN THE CLOUD AND ON THE GROUND

Michael T. Nygard  
Relevance, Inc.

[michael.nygard@thinkrelevance.com](mailto:michael.nygard@thinkrelevance.com) @mtnygard