BERLIN

SOFTWARE DEVELOPMENT



Training:Oct 15-16 / Conference:Oct 17-18

managing micro-services distributed systems and your infrastructure

jalewis@thoughtworks.com : @boicy : http://bovon.org



todays topics

recap on micro-services

some approaches to building them

some approaches to deploying them



Technology Radar

Prepared by the ThoughtWorks Technology Advisory Board

OCTOBER 2012

thoughtworks.com/radar

the story so far











hexagonal business capabilities owning their own data









Each capability decomposed into smaller applications based on your functional and cross-functional requirements



Sharing a uniform interface

200

DOD

1000





```
<html>
<body>
      <div>
          <span id="firstName">James</span>
          <span id="lastName">Lewis</span>
      </div>
      <div id="links">
          <a rel="self"
             href="/users/ae3fc"
             type="application/vnd.foobar.user+xml">Register User</a>
          <form id="update" method="POST" name="updateForm" action="/users/ae3fc">
              <input id="updatedFirstName" type="text" name="firstName" value="James"/>
              <input type="text" name="lastName" value="lewis"/>
              <input type="submit" name="updateButton" value="update"/>
          </form>
          <a rel="userIndex"
             href="/users/index"
             type="text/html">Search Users</a>>
      </div>
  </body>
</html>
```





Use standard application protocols to bridge the semantic gap

TIME FLIES LIKE AN ARROW



TIME LORDS LIKE A TARDIS





The **semantic** gap characterises the difference between two descriptions of an object by different linguistic representations, for instance languages or symbols.



The **semantic** gap characterises the difference between two descriptions of an object by different linguistic representations, for instance languages or symbols.



The **semantic** gap characterises the difference between two descriptions of an object by different linguistic representations, for instance languages or symbols.

RFC 5023 is an example









small, with a single responsibility





Object	Object
Object	Object



"objects should be no bigger than my head" my conjecture



and while I have a giant head, its pretty empty so thats ok...




Object Object	Object







as we chunk up **domains**, each **domain** should be small enough to fit in my head

in this case, it meant 100's of lines of code per application



Product Teams organised around Product Lines





testers



developers



THE BUSINESS







Conway's Law

Conway's Law

"...organizations which design systems ... are constrained to produce designs which are copies of the communication structure of those organizations"(sic)

Melvin Conway, 1968

Conway's Law

"...organizations which design systems ... are constrained to produce designs which are copies of the communication structure of those organizations"(sic)

Melvin Conway, 1968

"If you ask seven people to write a compiler, you get a seven pass compiler"

Dan North, circa 2006



independently deployable and scalable



























independently scalable, deployable, changeable, replaceable applications

whats the difference between building, deploying and testing this:



whats the difference between building, deploying and testing this:



and this?

~68% said yes

~ 32% said no

~80% are using automated provisioning

~20%

independently scalable, deployable, changeable, replaceable applications

independently scalable, deployable, changeable, replaceable applications

built, deployed and scaled automatically

building micro-services

I CONTRACTOR OF ADDRESS

hey, it works on my machine!



my machine!



hey, it works on I GUESS ITS TIME TO COMMIT AND PUSH



if you have lots of small applications all doing one thing

observing other resources, say



if you have lots of small applications all doing one thing

observing other resources, say



it is harder to put together a build pipeline for the components


Federated pipelines



Federated pipelines **Pre-release Production** testing **Provider** Consumer Provider pipeline runs contract tests

consumer tests to ensure the release won't break the consumer in production Consumer Sumer tests

Consumer team publishes test scripts to represent their requirements



there is a little bit more in here...



deploying many services

infrastructure automation



phoenix infrastructure

autoscaling and status aware





we saw earlier the basic build pipeline

we saw earlier the basic build pipeline



hey, it works on my machine!

> compile and integration performance unit test test test acceptance UAT deploy to production













its often only when you start pushing out to bigger environments that you start seeing problems



its often only when you start pushing out to bigger environments that you start seeing problems



and if the lovely Mr Nygard taught us anything it was to test with as realistic environments as early as possible





we can use virtualisation on our development machines to bring up whole environments at once



we can use virtualisation on our development machines to bring up whole environments at once

Vagrant



we can use virtualisation on our development machines to bring up whole environments at once

Vagrant

lxc

\$ vagrant box add base http://files.vagrantup.com/lucid32.box

- \$ vagrant init
- \$ vagrant up

\$ vagrant box add base http://files.vagrantup.com/lucid32.box

\$ vagrant init

\$ vagrant up

tough eh? Even I can use it

\$ vagrant box add base http://files.vagrantup.com/lucid32.box \$ vagrant init \$ vagrant up

tough eh? Even I can use it

under the hood, it applies a puppet manifest (or insert provisioning tool file format)

http://vagrantup.com/vI/docs/getting-started/index.html

multiple VM's are pretty easy too

multiple VM's are pretty easy too

```
Vagrant::Config.run do |config|
config.vm.define :web do |web_config|
web_config.vm.box = "web"
web_config.vm.forward_port 80, 8080
end
config.vm.define :db do |db_config|
db_config.vm.box = "db"
db_config.vm.forward_port 3306, 3306
end
end
```

multiple VM's are pretty easy too

```
Vagrant::Config.run do |config|
config.vm.define :web do |web_config|
web_config.vm.box = "web"
web_config.vm.forward_port 80, 8080
end
config.vm.define :db do |db_config|
db_config.vm.box = "db"
db_config.vm.forward_port 3306, 3306
end
end
```

it will take care of the networking business for you

http://vagrantup.com/v1/docs/getting-started/index.html

so now we can deploy our app locally to multiple VM's to test before pushing



so now we can deploy our app locally to multiple VM's to test before pushing



work on my machine!

so now we can deploy our app locally to multiple VM's to test before pushing



work on my machine!

which is a bit of a win in my book

we specify the dependencies for our application declaratively



we specify the dependencies for our application declaratively



```
class agent( $working_dir, $config_opts ) {
  require platform::sun-java
  require platform::endeavour-service
  File { owner => 0, group => 0, mode => 0644 }
  file { '/etc/motd':
    content => "Welcome to your Agent Node, managed by Puppet.\n"
  }
  platform::service::config {"service_configuration" :
    servicename => "agent",
    servicefilesource => "$working_dir/agent-1.0-SNAPSHOT.jar",
    servicefilename => "agent.jar",
    service_config_opts => $config_opts,
    require => Exec['install-java'];
  }
}
```

and we apply this to each environment we deploy our application to



```
class agent( $working_dir, $config_opts ) {
  require platform::sun-java
  require platform::endeavour-service
  File { owner => 0, group => 0, mode => 0644 }
  file { '/etc/motd':
    content => "Welcome to your Agent Node, managed by Puppet.\n"
  }
  platform::service::config {"service_configuration" :
    servicename => "agent",
    servicefilesource => "$working_dir/agent-1.0-SNAPSHOT.jar",
    servicefilename => "agent.jar",
    service_config_opts => $config_opts,
    require => Exec['install-java'];
  }
}
```
and we apply this to each environment we deploy our application to



class agent{ Sworking_dir, Sconfig_opts } { require platform::sum-java require platform::endeavour-service File { owner => 0, group => 0, mode => 0644 }

content => "Welcome to your Agent Node, managed by Puppet.\n"
}
platform::service:config ("service_configuration" :
servicements => "gener":
servicefilesource => "generking_dir/agent-1.0-BNAPSHOT.jar",
servicefilesome => const.iar".

servicefilesource => "\$working_dir/agent-1.0-SNAPSHOT. servicefilename => "agent.jar", service_config_opts => \$config_opts, require => Exec['install-java'];

and we apply this to each environment we deploy our application to



and we apply this to each environment we deploy our application to



and we **abstract** away the fact that we are pushing to a local Data Centre or to an laaS platform like AWS









and there is tooling available to do this too

and there is tooling available to do this too

```
@task
```

```
def deploy_environment(environment, build_number):
    target_env = ENVIRONMENTS[environment]
    deployment = Deployment(build_number,target_env)
    deployment.install(SERVICES)
```

@task

```
def end_to_end_test_against(environment, build_number):
    target_env = ENVIRONMENTS[environment]
    deployment = Deployment(build_number,target_env)
    deployment.run_end_to_end_tests(EndToEndTest())
```

and there is tooling available to do this too

```
@task
def deploy_environment(environment, build_number):
    target_env = ENVIRONMENTS[environment]
    deployment = Deployment(build_number,target_env)
    deployment.install(SERVICES)

@task
def end_to_end_test_against(environment, build_number):
    target_env = ENVIRONMENTS[environment]
    deployment = Deployment(build_number,target_env)
    deployment.run end to end tests(EndToEndTest())
```

this is a tool called Fabric

it's a wrapper over SSH so you can run commands against anything you have SSH access to

it's a wrapper over SSH so you can run commands against anything you have SSH access to

```
class RemoteServer(ApplicationServer):
    def deploy_service(self, target_env, service_name, artefact_repo):
        with settings(user='service', host_string=self.host, key_filename='~/.ssh/
deploy'):
        run('rm -rf /tmp/deploy.%s' % artefact_repo.build_number)
        run('svn checkout /trunk/platform_deploy /tmp/deploy.%s' %
            artefact_repo.build_number
        )
        run('/tmp/deploy.%s/deploy-no-db.sh %s %s %s %s' % (
            artefact_repo.build_number,
            service_name,
            artefact_repo.build_number,
            self.install_dir,
            target_env.property_file_name)
        )
    }
```

```
class LocalServer(ApplicationServer):
    def deploy_service(self, target_env, service_name, artefact_repo):
        local('rm -rf /tmp/deploy.%s' % artefact_repo.build_number)
        local('svn checkout /trunk/platform_deploy /tmp/deploy.%s' %
            artefact_repo.build_number
        )
        local('/tmp/deploy.%s/deploy-no-db.sh %s %s %s %s %s' % (
            artefact_repo.build_number,
            service_name,
            artefact_repo.build_number,
            service_name,
            artefact_repo.build_number,
            self.install_dir,
            target_env.property_file_name)
```

```
class RemoteServer(ApplicationServer):
    def deploy_service(self, target_env, service_name, artefact_repo):
        with settings(user='service', host_string=self.host, key_filename='~/.ssh/
deploy'):
    run('rm -rf /tmp/deploy.%s' % artefact_repo.build_number)
    run('svn checkout /trunk/platform_deploy /tmp/deploy.%s' %
        artefact_repo.build_number
    )
    run('/tmp/deploy.%s/deploy-no-db.sh %s %s %s %s %s' % (
        artefact_repo.build_number,
        service_name,
        artefact_repo.build_number,
        self.install_dir,
        target_env.property_file_name)
```

and yes, the astute among you will have noticed the awful duplication

```
class LocalServer(ApplicationServer):
    def deploy_service(self, target_env, service_name, artefact_repo):
        local('rm -rf /tmp/deploy.%s' % artefact_repo.build_number)
        local('svn checkout /trunk/platform_deploy /tmp/deploy.%s' %
            artefact_repo.build_number
        )
        local('/tmp/deploy.%s/deploy-no-db.sh %s %s %s %s %s' % (
            artefact_repo.build_number,
            service_name,
            artefact_repo.build_number,
            service_name,
            artefact_repo.build_number,
            self.install_dir,
            target_env.property_file_name)
        )
    }
}
```

```
class RemoteServer(ApplicationServer):
    def deploy_service(self, target_env, service_name, artefact_repo):
        with settings(user='service', host_string=self.host, key_filename='~/.ssh/
deploy'):
    run('rm -rf /tmp/deploy.%s' % artefact_repo.build_number)
    run('svn checkout /trunk/platform_deploy /tmp/deploy.%s' %
        artefact_repo.build_number
    )
    run('/tmp/deploy.%s/deploy-no-db.sh %s %s %s %s %s %s % (
        artefact_repo.build_number,
        service_name,
        artefact_repo.build_number,
        self.install_dir,
        target_env.property_file_name)
```









declarative environment provisioning

```
prod:
  nodes:
  - ami id: ami-4dad7424
    size:
          t1.micro
    credentials name: us-east-ssh
    aws key name : test
    services: [hello world]
    security groups: [ spicy-beef ]
    availability zone: us-east-1a
    type: phoenix.providers.aws provider.AWSNodeDefinition
  - ami id: ami-4dad7424
    size: t1.micro
    credentials name: us-east-ssh
    aws key name : test
    services: [hello world]
    security groups: [ spicy-beef ]
    availability zone: us-east-1b
    type: phoenix.providers.aws provider.AWSNodeDefinition
  - ami id: ami-4dad7424
    size: t1.micro
    credentials name: us-east-ssh
    aws key name : test
    services: [apache]
    type: phoenix.providers.aws provider.AWSNodeDefinition
    security groups: [ spicy-beef ]
  node provider:
    class name: AWSNodeProvider
    public api key: {{ aws public api key }}
    private api key: {{ aws private api key }}
```

declarative environment provisioning

prod:			
nodes:			
<pre>- ami_id: ami-4dad7424</pre>			
size: t1.micro	apache:		
credentials_name: us-east	puppet module directory : puppet		
aws_key_name : test	puppet manifest : apache.pp		
<pre>services: [hello_world]</pre>	service configurator:		
security_groups: [spicy-]	phoenix.configurators.puppet service configurator.PuppetServiceConfigurator		
availability_zone: us-eas	connectivity:		
type: phoenix.providers.av	- protocol: tcp		
- ami_id: ami-4dad7424	ports: [80]		
size: t1.micro	allowed: [WORLD]		
credentials_name: us-east			
aws_key_name : test	hello_world:		
<pre>services: [hello_world]</pre>	puppet_module_directory : puppet		
<pre>security_groups: [spicy_]</pre>	- puppet manifest : hello world.pp		
availability_zone: us-eas	service_configurator:		
type: phoenix.providers.av	phoenix.configurators.puppet_service_configurator.PuppetServiceConfigurator		
- ami_id: ami-4dad7424	connectivity:		
size: tl.micro	- protocol: tcp		
credentials_name: us-east	ports: [8080, 8081]		
aws_key_name : test	allowed: [WORLD]		
services: [apache]			
type: phoenix.providers.av	mongo:		
security_groups: [spicy-]	<pre>puppet_module_directory : puppet</pre>		
	<pre>puppet_manifest : mongo.pp</pre>		
node_provider:	service_configurator:		
class_name: AWSNodeProvide	phoenix.configurators.puppet_service_configurator.PuppetServiceConfigurator		
<pre>public_api_key: {{ aws_pub</pre>	connectivity:		
private_api_key: {{ aws_p	- protocol: tcp		
	ports: [27017]		
	allowed: [hello_world]		

declarative environment provisioning

<pre>prod: nodes: - ami id: ami-4dad7424</pre>		
size: t1.micro credentials_name: us-east- aws_key_name : test service: security availab	apache: -sspuppet_module_directory : puppet puppet_manifest : apache.pp rator: tors.puppet_service_cor	figurator.PuppetServiceConfigurator
type: pl	VUINS pefinition	
size: t1.micro credentials name: us-east-	allowed: [WORLD] -ssh	
aws_key_name : test services: [hello_world] security_groups: [spicy-k	hello_world: puppet_module directory : puppet peepuppet manife	
availability_zone: us-east type: phoenix.providers.av - ami_id: ami-4dad7424	connectivity:	Tate ppetServiceConfigurator
<pre>size: t1.micro credentials_name: us-east- aws_key_name : test</pre>	<pre>- protocol: -ssh ports: [8080, 8081] allowed: [WORLD]</pre>	
services: [apache] type: phoenix.providers.av	vmongovider.AWSNodeDefinition	
security_groups: [spicy-1	peefuppet_module_directory : puppet qqqqqqqp	
node provider:	service configurator:	
class name: AWSNodeProvide	The service configurators number service con	figurator.PuppetServiceConfigurator
public api kev: {{ aws pub	liconpickevity.	
private api key: {{ aws pi	civate apatkew]}tcp	
	ports: [27017]	
	allowed: [hello_world]	

Deployment



the Fowler bomb

the Fowler bomb

"it pushes the accidental complexity into the infrastructure"

Martin Fowler

the Fowler bomb

"it pushes the accidental complexity into the infrastructure"

Martin Fowler

turns out that monitoring and logging have a big part to play too





each of these should report metrics at a well known location



each of these should report metrics at a well known location

these metrics can displayed locally



each of these should report metrics at a well known location

these metrics can displayed locally

and pushed to centralised monitoring tools



each of these should report metrics at a well known location

these metrics can displayed locally

and pushed to centralised monitoring tools

Status aware applications are the way forward

graphite for example


















the application could block incoming requests, or notify operations

in java land there is a nice library from the good folks at yammer that allows you to do a lot of this

```
import com.yammer.metrics.core.HealthCheck;
public class NamedHealthCheck extends HealthCheck {
    private final Pingable pingable;
    private final String name;
    public NamedHealthCheck(Pingable pingable, String name) {
        super(name);
        this.pingable = pingable;
        this.name = name;
    }
    @Override
    protected Result check() {
        try {
            return (pingable.ping())
                       ? Result.healthy()
                       : Result.unhealthy(String.format("Could not connect to the %s", name));
        } catch (Exception e) {
            return Result.unhealthy(String.format("Could not connect to the %s", name));
        }
    }
```

http://metrics.codahale.com

then there are the groovy new (and some not so new) tools on the block

then there are the groovy new (and some not so new) tools on the block

- Graylog2
- Syslog-ng
- Logstash
- Scribe
- Splunk
- riemann

then there are the groovy new (and some not so new) tools on the block

- Graylog2
- Syslog-ng
- Logstash
- Scribe
- Splunk
- riemann

and a final shout out to zipkin for distributed system monitoring





but "invented a slightly better one. That finally got changed once more to what we have today. He put pipes into Unix." Thompson also had to change most of the programs, because up until that time, they couldn't take standard input. There wasn't really a need; they all had file arguments. "GREP had a file argument, CAT had a file argument."

The next morning, "we had this orgy of `one liners.' Everybody had a one liner. Look at this, look at that. ...Everybody started putting forth the UNIX philosophy. <u>Write</u> programs that do one thing and do it well. Write programs to work together. Write programs that handle text streams, because that is a universal interface." Those ideas which add up to the tool approach, were there in some unformed way before pipes, but they really came together afterwards. Pipes became the catalyst for this UNIX philosophy. "The tool thing has turned out to be actually successful. <u>With pipes, many</u> programs could work together, and they could work together at a distance."

The Unix Philosophy

Lions commentary on Unix 2nd edition

to build systems is to make trade-offs

to build systems is to make trade-offs

maintainability vs time-to-market

throughput vs cost

portability vs deployability

to build systems is to make trade-offs

maintainability vs time-to-market

throughput vs cost

portability vs deployability

infrastructure automation and tooling is essential to get the benefits of this approach

BERLIN

SOFTWARE DEVELOPMENT



Training:Oct 15-16 / Conference:Oct 17-18

thanks

jalewis@thoughtworks.com : @boicy : http://bovon.org

ThoughtWorks*

BERLIN

SOFTWARE DEVELOPMENT CONFERENCE 2013



Training:Oct 15-16 / Conference:Oct 17-18

ThoughtWorks[®] is hiring

jalewis@thoughtworks.com : @boicy : http://bovon.org