

Changing The Engine While In Flight

Neil Armitage
Senior DevOps Engineer, VMware



 **puppetconf**
2016

vmware[®]

© 2014 VMware Inc. All rights reserved.

WHOAMI

- Senior DevOps Engineer at VMware focusing on internal cloud deployments.
- Nearly 30 Years of Ops/DBA/Developer experience from IBM Mainframes upwards.
- Based in Palo Alto but work from the Scottish Highlands 😊



Background


- In Oct 2014 VMware acquired the assets of Continuent Inc
- The Continuent team joined VMware's Hybrid Cloud Business Unit.
- Focusing on bringing DBaaS into vCloudAir
- Needed to migrate Continuent Test/Dev/QA Systems from a mix of outsourced resources into a new internal vSphere Cluster
- Needs to be non disruptive as product launches are planned

Note
There were some last minute
changes for Continuent which made
some of the slides inconsistent

What is(was) Continuent

- Commercial Continuent Tungsten focused on MySQL asynchronous Clustering
- Open Source Tungsten Replicator, moving data from
 - MySQL to MySQL
 - Oracle to MySQL
 - MySQL to Oracle
 - MySQL or Oracle to Hadoop
 - MySQL or Oracle to Redshift
- Around 20 globally dispersed Engineers and Support staff

Where were our servers?



What we had

- Around 50 Physical and Virtual Linux Hosts running
 - Customer facing Website (Joomla)
 - Jenkins Environment
 - Test and QA Clusters
 - Support Jump hosts for accessing Customer sites
 - Puppet Master
- All with different configurations some going back 10 years
- Some under Puppet control mainly covering Users and Firewalls
- Centos 4,5 & 6, Ubuntu 12.14.....

What we had

- CI Pipelines in Jenkins containing
 - 10+ build jobs
 - 200+ Unit and Integration tests
 - Integration tests running against MySQL, Oracle, Hadoop and AWS Redshift.

Why Puppet

- A few years ago we compare Puppet vs Chef and getting started with Puppet was easier
- Looked at Ansible when it matured but didn't see it as a good choice, centralized server made sense
- Not a Puppet 'fanboy', it's a tool in our toolbox.

State Pre migration


- Several machines already 'puppetized'
- Initial adoption was triggered by several hacks so the modules concentrated on
 - Firewalls – controlling ingress into the nodes
 - Users – disabling root, maintaining SSH keys for users
 - Moving SSH to a new port
 - Using a jump host as a gateway
 - Initial separate puppet module for tungsten setup (now forked into a OSS module)
 - <https://github.com/continuent/continuent-puppet-tungsten>

Why Migrate


- VMware not keen on paying AWS 😊
- Dealing with multiple vendors was hard
- Hardware was old and no longer met our requirements
- We had around 40 QA hosts QA Team wanted 400+
- Move from external Subversion to internal Git

Where we were going

- Brand new vSphere 6 cluster running vSAN - 29 x Dell PE R730xd; 24C, 512GB
- Around 300TB of shared vSAN Disk
- 70 x Dell PE R730xd; 12C, 128GB for physical host testing (Hadoop etc)
- Totally isolated only port 80 and 443 to outside world.



What is vSAN?




Constraints/Concerns

- We had committed to ship multiple releases of Continuent Tungsten post acquisition
- We had to ship them as customers needed re-assurance
- We couldn't break the QA environment based on 1 and 2 above
- The environment we were moving into was new and we had limited vCenter knowledge



Plan



Existing QA Environment

Build new Environment

Parallel Run

New Enviroment (Take 1)

- 29 Hosts Clustered into a single vCenter environment
- Single vSAN Cluster of 320TB
- Deployed a Puppet Master and PuppetDB server
- Started work on new modules

2 days later

- All the VM's deployed had gone
- vSAN cluster had failed
- It appears some one had purchased SSD's which were not supported by vSAN
- (this took about 2 weeks to discover)



New environment (Take 2)

- 29 Hosts Clustered into a single vCenter environment
- ESX hosts set up to use both local disks and a borrowed VNX San
- Deployed a Puppet Master and PuppetDB server
- Started work on new modules


2 days later



Infrastructure

Manual

Puppet



Physical Network

Virtual Hosts

Virtual Network

Puppet modules


- 'Base' class applied to all hosts
 - Users and SSH keys
 - Default packages per O/S – Centos and Ubuntu initially
 - Remote syslog
 - NTP
 - Nagios
 - eth1 Management
- RDBMS specific
- Jenkins and Monitoring rely heavily on exported resources from the Base class.

What are exported resources?

“An exported resource declaration specifies a desired state for a resource, does not manage the resource on the target system, and publishes the resource for use by other nodes. Any node (including the node that exported it) can then collect the exported resource and manage its own copy of it.”

https://docs.puppet.com/puppet/latest/reference/lang_exported.html

What are exported resources?



DNS Server Management with Exported Resources

```
export_a_record.pp  x
1  #Record info in PuppetDB for the DNS Server
2  @@dns::record::a { $hostname:
3      zone => $::domain,
4      data => $::ipaddress_eth0,
5  }
6
```

DNS Server Management with Exported Resources

```
import_a_record.pp  x
1 node /^ns\d+\.wdc3\.continuent\.com$/ {
2
3   include dns::server
4
5   dns::server::options { '/etc/bind/named.conf.options':
6     forwarders => [ '10.142.7.1', '10.142.7.2' ]
7   }
8
9   dns::zone { 'wdc3.continuent.com':
10     soa      => 'ns1.wdc3.continuent.com',
11     ...
12   }
13
14   #Pull all of the A records from PuppetDB
15   Dns::Record::A <<| |>>
16 }
17
```

Demo

Demo Setup



PuppetMaster



WEB1



MYSQL1



DNS1

Demo Setup

```
demo_setup.pp  x
1  yum install puppetserver
2  puppet module install puppetlabs-puppetdb
3
4  Site.pp
5
6  node puppetserver {
7      # Configure puppetdb and its underlying database
8      class { 'puppetdb': }
9      # Configure the Puppet master to use puppetdb
10     class { 'puppetdb::master::config': }
11 }
```


QA Cluster

- Built in groups of 3,6,9 or 12 nodes
- QA Class
- Added RDBMS as specified
- Extra QA tools, debugging etc.

MySQL QA Cluster

```
mysql_qa.pp  x
1 node /^c6-az-6-\d+\.wdc3-qa\.com$/ {
2   class {"continuent_qa_node_2":
3     nodes_data => {
4       host1=> {'host'=>'c6-01', 'iIp' => '192.168.12.216', 'eIp' => '10.158.12.216' },
5       host2=> {'host'=>'c6-02', 'iIp' => '192.168.12.217', 'eIp' => '10.158.12.217' },
6       host3=> {'host'=>'c6-03', 'iIp' => '192.168.12.218', 'eIp' => '10.158.12.218' },
7     } ,
8     mysqlVersion=>'ps5.5.32'
9   }
10 }
```

Oracle QA Cluster

```
oracle_qa.pp
1 node /^o11-bf-2-\d+\.wdc3-qa\.com$/ {
2   class {"continuent_qa_node_2":
3     nodes_data => {
4       host1=>{'host'=> 'o11-01','iIp'=>'192.168.12.242','eIp'=> '10.158.12.242' },
5       host2=>{'host'=> 'o11-02','iIp'=>'192.168.12.243','eIp'=> '10.158.12.243' },
6     } ,
7     installMySQL=>true,
8     installOracle=>true,
9     oracleVersion=>11,
10    nfsMount=>"10.158.12.27:/nfs/oracle"
11  }
12 }
```

RDBMS Supported

- MySQL
 - Oracle MySQL
 - MariaDB
 - Percona Server
- Oracle EE – 11g and 12c
- Vertica
- Hadoop

Jenkins configuration

- Several hundred tests in Jenkins
- Pre-migration each test specified a cluster to run on.
- Led to bottle necks and problems when a cluster is unavailable
- In the new environment a test just specifies the number of nodes and the O/S it needs

☒ Execute concurrent builds if necessary [?](#)

☒ Restrict where this project can be run [?](#)

Label Expression [?](#)

[Label](#) is serviced by 8 nodes

Advanced Project Options

Advanced...

Jenkins configuration

- Puppet creates the Jenkins slave using data from exported resources.
- Metadata inserted into the workspace by puppet to allow the test to find the correct hosts



Slave c6-am-6 (c6-am-6)

Created by [Neil Armitage](#)

Labels

[c6-1](#) [c6-2](#) [c6-3](#) [c6-4](#) [c6-5](#) [c6-6](#)

Projects tied to c6-am-6

None



Jenkins configuration

```
hosts_with_nodes.sh x
1 $ cat /var/lib/jenkins/exec/c6-am-6/hosts_with_node.txt
2 NODE1=c6-am-6-01.wdc3-qa.com
3 NODE2=c6-am-6-02.wdc3-qa.com
4 NODE3=c6-am-6-03.wdc3-qa.com
5 NODE4=c6-am-6-04.wdc3-qa.com
6 NODE5=c6-am-6-05.wdc3-qa.com
7 NODE6=c6-am-6-06.wdc3-qa.com
8
```

Completed Environment

- VM's deployed using PowerShell to clone template, set hostname and add IP for eth0
- Nodes booted and ran puppet
- Internal DNS was set correctly in template so puppet agent found the puppetmaster
- Node configured from puppet master
- Monitoring automatically populated on Nagios hosts when puppet ran on that host
- DNS records updated in DNS servers
- Cluster registered itself with Jenkins server as a new available node via exported resources

Parallel Running

- Tests manually copied from old Jenkins host to new host
- Tests ran in parallel for approx. 1 month
- The only real difference was run time 1day on old env -> 1 hour in new env
- Old environment was decommissioned

Enhancements

- Needed to start using Windows and SQL Server
- Played with puppet enterprise to look at the puppet sql server module
- Could see the use but it took too long to get the PO approved.

Future

- VMware EOL all Continuent products in May 2016.
- Continuent Software is being spun back off into a separate Company.
- Currently working on migrating environment back to AWS (using Puppet).
- About 75% of the environment has now been decommissioned and reallocated to new projects.
- Lessons learnt have been carried through to the next project.

Lessons Learnt

- Initial investment is high but the long term payoff is good
- Resist the temptation to go a quick hack rather than modify the puppet module
- Resist the temptation to go a quick hack rather than modify the puppet module
- We had lots of issues around memory usage on puppetDB when running 3.7.x
 - Allocate lots of JVM memory
 - Not run 4.0.x at the same scale yet so I don't know if it's fixed.
- Make sure modules are in a SCS system we use Git.
 - Develop locally and push to a repo
 - Puppet Master pulls the latest code

Questions



Slides and Code
<https://github.com/narmitag/puppetconf2016>