

# Setting up an HPC lab from scratch

with Mr-Provisioner, Jenkins, OpenHPC and Ansible

FOSDEM 2019 - Feb 3rd - HPC devroom  
Baptiste Gerondeau  
Renato Golin



# Setting Up

- Lab bring up
  - KVM (libvirt) bringup: provisioner, jenkins, file server
  - Network setup: static IPs, split BMC / provisioning networks
  - Ansible used to install Jenkins, provisioner, SFTP, apt-cache
  - Private Git repository
- Mr-Provisioner: <https://github.com/mr-provisioner/mr-provisioner>
  - Bare metal provisioning: BMC control, PXE, boot, console
  - DHCP/TFTP service: dynamic images per MAC address (via Grub scripts)
  - Machine discovery, user management, preseed/image upload, JSON automation
- Jenkins: [https://github.com/Linaro/hpc\\_lab\\_setup](https://github.com/Linaro/hpc_lab_setup)
  - Standard install with LDAP + Jenkins Job Builder
  - Install / update jobs via Ansible
  - Jobs checkout same repo for scripts & dynamic execution
- File Server
  - SFTP for benchmark / jobs results (user control), apt-cache, temp web for toolchains

# Provisioning



## qdfohpc

Overview

Power state

on

BMC

[bmc-qdf-openhpc](#)

Architecture

[AArch64](#)

Interfaces +

Identifier ↓	MAC	Network	Configuration type	Configured IPv4	Lease IPv4	Last seen		
BMC	70: <input type="text"/>	bmc_network (10.41.0.0/16)	dynamic-reserved	10.41.24.120	10.41.24.120	30 Jan 2019 09:36	<a href="#">Edit</a>	<a href="#">Remove</a>
eth1	8c: <input type="text"/>	machine_network (10.40.16.0/20)	dynamic-reserved	10.40.24.120	10.40.24.120	30 Jan 2019 09:21	<a href="#">Edit</a>	<a href="#">Remove</a>

### Actions

- Console
- Reset console
- PXE Reboot
- Reboot
- Power off
- Delete Machine

# Provisioning



## Provisioning

Subarchitecture

GrubWithOptionEfiboot

Kernel

bger/vmlinuz1.d058e26c (CentOS 7.5)

Kernel

ip=dhcp text

cmdline

inst.stage2=http://10.40.0.13/mirror.centos.org/altarch/7/os/aarch64/  
inst.repo=http://10.40.0.13/mirror.centos.org/altarch/7/os/aarch64/  
inst.ks=file:/ks.cfg earlycon console=ttyAMA0,115200

Initrd

bger/initrd1.img.f7025fca (CentOS 7.5)

Preseed

CentOS.Upstream (CentOS Upsteam)

Netboot enabled

☒

## Assignees +

User ↓	Reason	Assignment date	
bger		19 Apr 2018 20:11	<a href="#">Edit</a> <a href="#">Remove</a>
jenkins		20 Jun 2018 11:26	<a href="#">Edit</a> <a href="#">Remove</a>
rengolin		19 Apr 2018 20:11	<a href="#">Edit</a> <a href="#">Remove</a>

## Events

Show all ☐

 Console accessed by <b>bger</b>	29 Jan 2019 16:38:56
 DHCPREQUEST packet seen	29 Jan 2019 16:35:10
 DHCPREQUEST packet seen	29 Jan 2019 16:20:36

# Provisioning



Console for **qdf0hpc**













```
[ OK ] Stopped target Initrd Root File System.
[ OK ] Listening on udev Control Socket.
       Mounting Debug File System...
       Starting Collect Read-Ahead Data...
       Mounting Huge Pages File System...
[ OK ] Stopped target Initrd File Systems.
[ OK ] Reached target Slices.
[ OK ] Listening on /dev/initctl Compatibility Named Pipe.
[ OK ] Created slice system-serial\x2dgetty.slice.
[ OK ] Created slice system-selinux\x2dpol...grate\x2dlocal\x2dchanges.slice.
[ OK ] Started Collect Read-Ahead Data.
       Starting Remount Root and Kernel File Systems...
       Starting Apply Kernel Variables...
[ OK ] Started Apply Kernel Variables.
[ OK ] Started Create list of required sta...ce nodes for the current kernel.
       Starting Create Static Device Nodes in /dev...
[ OK ] Started Journal Service.
[ OK ] Mounted POSIX Message Queue File System.
[ OK ] Mounted Huge Pages File System.
[ OK ] Mounted Debug File System.
[ 12.112922] EXT4-fs (sda3): re-mounted. Opts: (null)
[ OK ] Started Read and set NIS domainname from /etc/sysconfig/network.
[ OK ] Started Remount Root and Kernel File Systems.
       Starting Rebuild Hardware Database...
```

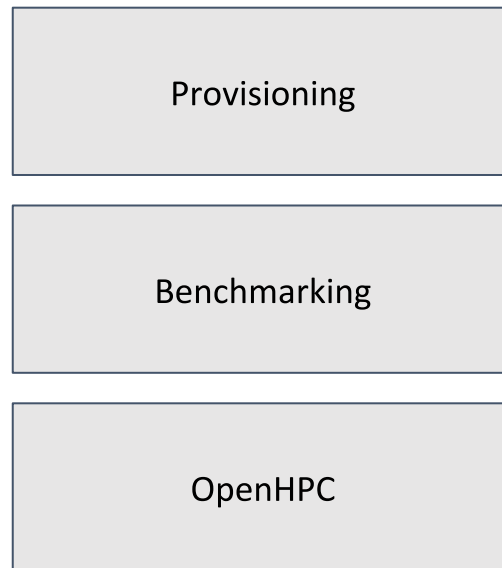
Actions

 Download log

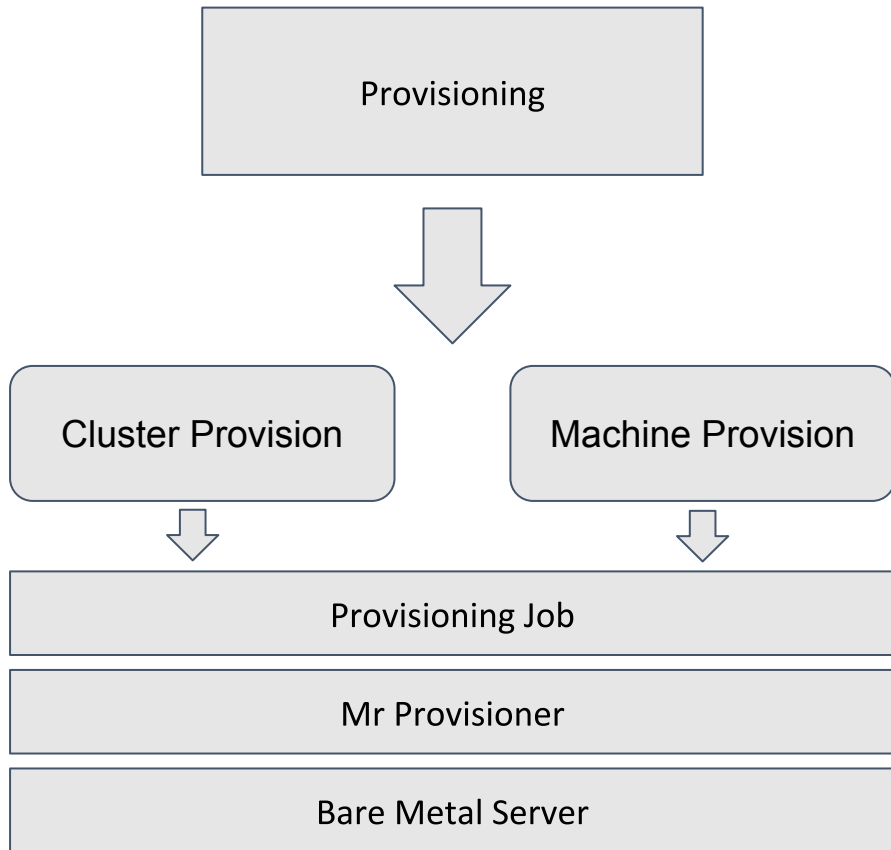
# Jenkins Jobs

All	Benchmarks	OpenHPC	Provisioning	+
-----	------------	---------	--------------	---

S	W	Name ↓
		benchmark_job
		cluster_provision
		machine_provision
		ohpc_install
		ohpc_test_suite
		provisioning_job



# Jenkins Jobs



## Project cluster\_provision

This build requires parameters:

machine_type	<input type="text" value="d05"/>
The cluster of machines to be provisioned	
client_branch	<input type="text" value="master"/>
The branch of the mrp client to use	
automation_branch	<input type="text" value="master"/>
The Ansible logic branch to use	

Build

## Project machine\_provision


This build requires parameters:

machine_type	<input type="text" value="d03"/>
Micro-architecture of the machine to be provisioned	
job_type	<input type="text" value="bench"/>
The job type to be executed on the machine	
os_type	<input type="text" value="default"/>
The OS distribution. Default: Debian for benchmark, CentOS for openhpc	
client_branch	<input type="text" value="master"/>
The branch of the MrP client to use	
automation_branch	<input type="text" value="master"/>
The Ansible logic branch to use	

Build

# Jenkins Jobs

Benchmarking



Benchmark Job



benchmark\_harness

Benchmarked Machine

## Project benchmark\_job


This build requires parameters:

node	<div>d03bench d05bench qdfbench tx2bench</div>	Node with which to run the benchmark on the machine
branch	<div>master</div>	The benchmark_harness branch to use
automation_branch	<div>master</div>	The Ansible logic branch to use
benchmark	<div>lulesh</div>	The benchmark to run on the machine
size	<div>3</div>	The benchmark specific size parameter. 1 is small, 3 is large.
iterations	<div>1</div>	The number of times to run the binaries (statistical purposes)
compiler	<div>gcc</div>	The compiler (url or system binary) with which to compile the benchmark
compiler_flags	<div></div>	The extra compiler flags to compiler with
linker_flags	<div></div>	The extra linker flags to link with
run_flags	<div></div>	The extra execution flags to run the benchmark
harness_options	<div></div>	Additional harness options

Build



# Jenkins Jobs



## Project ohpc\_install

This build requires parameters:

node	<div>d05ohpc qdfohpc tx2ohpc</div>
Node on which to execute this job	
method	<div>stateful</div>
The type of OHPC install to do	
client_branch	<div>master</div>
Branch name of the mr-provisioner-client to use	
automation_branch	<div>linaro</div>
Branch name of the ohpc install recipe	

Build


## Project ohpc\_test\_suite

This build requires parameters:

node	<div>d05ohpc qdfohpc tx2ohpc</div>
Node on which to execute this job	
method	<div>stateful</div>
The type of OHPC install to do	
test_type	<div>long</div>
The length of the tests	
client_branch	<div>master</div>
Branch name of the ohpc install recipe	
automation_branch	<div>linaro</div>
Branch name of the ohpc install recipe	

Build


# Jenkins Jobs


 **Jenkins**


2 search ? BAPTISTE GERONDEAU | LOG OUT


Jenkins > ohpc\_test\_suite > DISABLE AUTO REFRESH


[Back to Dashboard](#)


 **Status**


 Changes


 Workspace

 Build with Parameters

 Delete Project


 Configure


 Rebuild Last


 Rename


## Project ohpc\_test\_suite

This is the job to run the OpenHPC test suite on a cluster<!-- Managed by Jenkins Job Builder -->

 Workspace









 Recent Changes

 Latest Test Result (no failures)

 **Build History**

trend ^

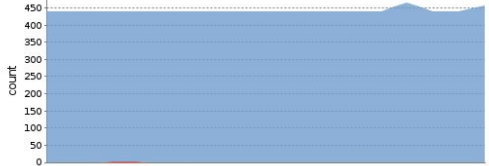
find X

 <b>#65</b>	29-Jan-2019 13:38	
 <b>#64</b>	29-Jan-2019 12:13	
 <b>#63</b>	24-Jan-2019 11:21	
 <b>#62</b>	11-Jan-2019 14:52	


### Permalinks

- Last build (#65), 3 hr 4 min ago
- Last stable build (#65), 3 hr 4 min ago
- Last successful build (#65), 3 hr 4 min ago
- Last failed build (#64), 4 hr 29 min ago
- Last unstable build (#41), 2 mo 7 days ago
- Last unsuccessful build (#64), 4 hr 29 min ago
- Last completed build (#65), 3 hr 4 min ago

### Test Result Trend

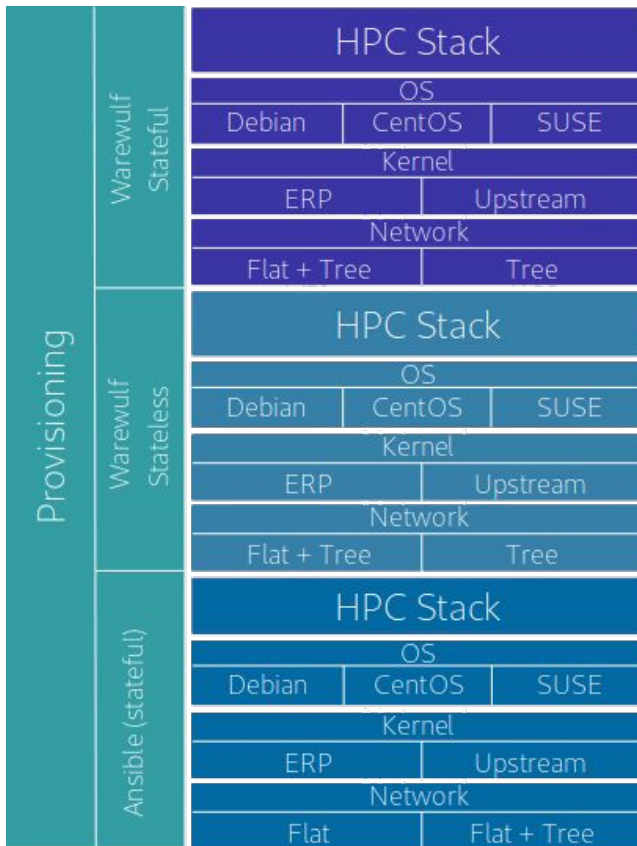


(just show failures) enlarge

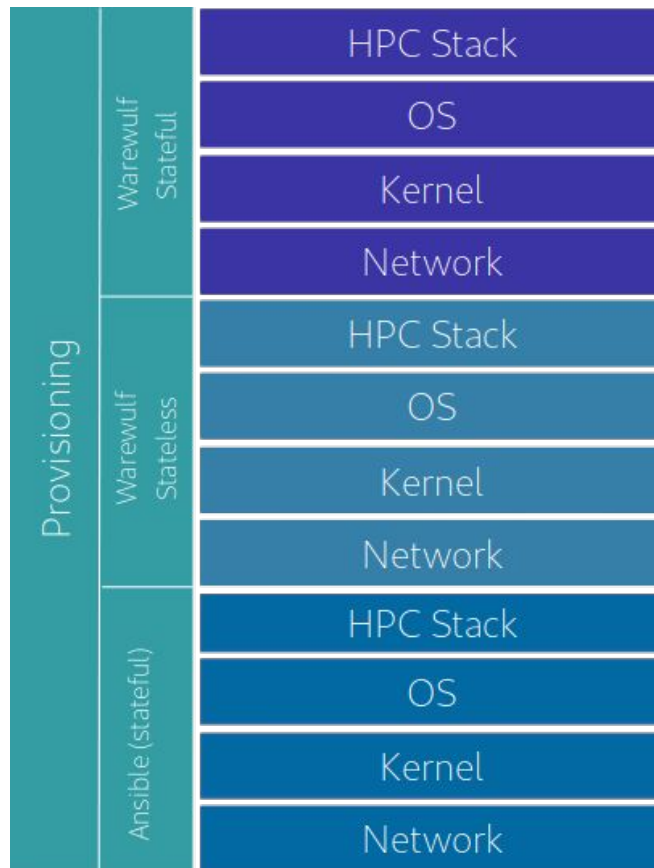
 edit description

**Disable Project**

# OpenHPC



# OpenHPC



# Conclusion

- Modularity
  - Iterative development leading to modular design
- Multiple configurations for each machine/cluster
  - Support for OpenHPC installation on OpenSUSE coming very soon
- Low Maintenance
  - Automation, versioning for the SDI and Jenkins Jobs
- Jenkins Job Builder, Ansible and Python
  - Well known, well used, accessible technologies
- Multiple provisioning options in OpenHPC
  - Ansible stateful only a “bonus” to accomodate for (any) network structure
  - Ansible easily runnable in degraded mode (ignoring certain machines in the cluster)
  - Warewulf Stateless coming very soon
  - Warewulf Stateful coming

Thanks!

