Working with Puppet-managed physical machines

Vítor Gouveia,
vitor.gouveia@cern.ch

IT-PES-PS

Agenda

- Puppet-managed physical machines
  - quattor-managed physical machines
  - partition table
  - Kickstart provisioning templates
  - Installing a physical machine
- PuppetDB

Puppet-managed physical machines

Physical machines will only be allocated in the event that the requirements can not be meet with virtual machines

Puppet-managed physical machines

Quattor-managed physical machines:

- cannot “convert” your physical machines
- you need always to request a new puppet-managed physical machine and release the Quattor-managed machines that you do not need anymore: KB0002310 – how to release quattor managed machines
  - Why?
    - to ensure that all the inconsistencies between hardware/network/configuration DBs are ironed out before hardware is re-allocated
    - to ensure that the hardware supports SLC6/Puppet
    - to ensure that the Foreman can be populated with the correct information regarding IPMI, BMC etc.

Puppet-managed physical machines

Request:

- HW allocation request form
  - Attention: You cannot choose the machine name due to the host naming policy: KB0002301

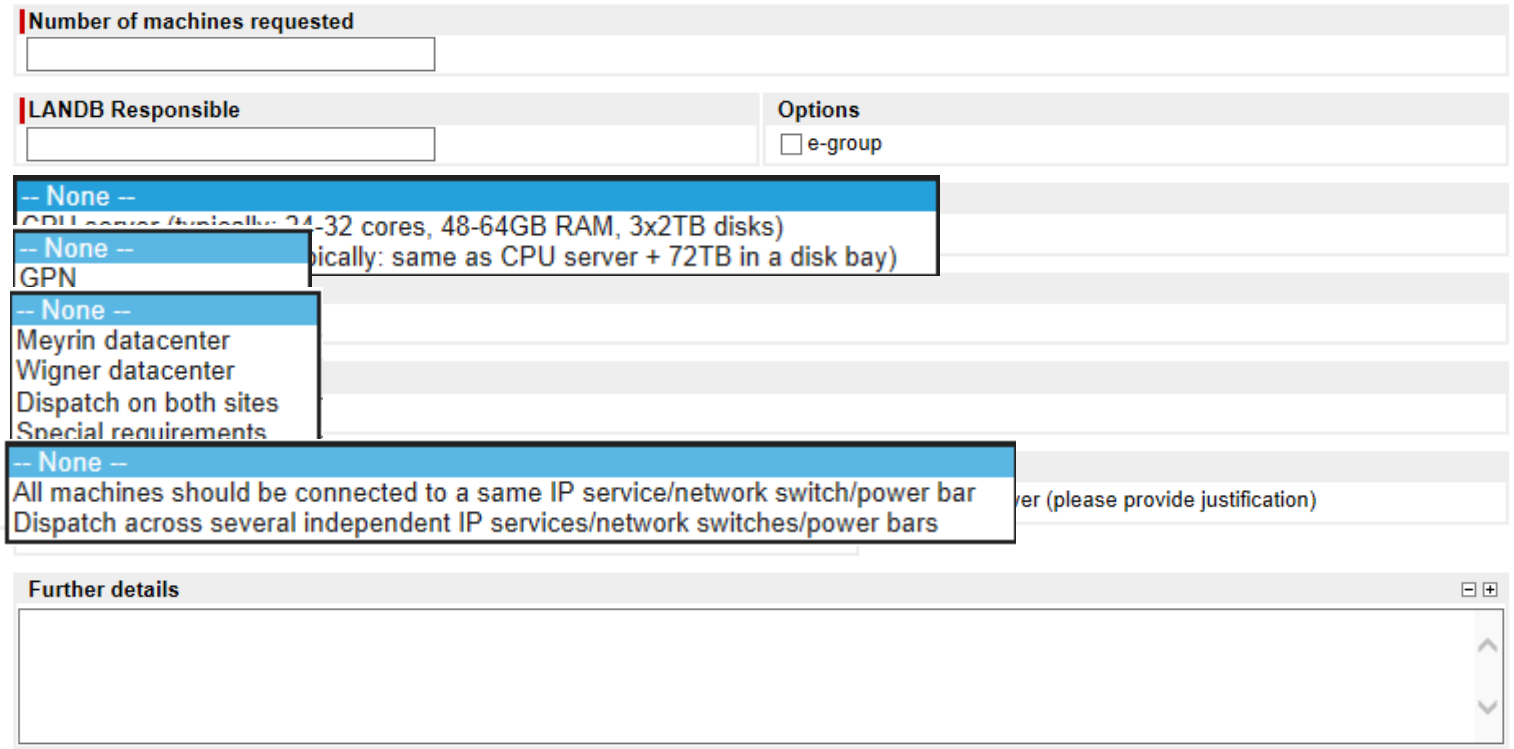

Puppet-managed physical machines

Why you cannot choose the name of physical machine…

1. new hosts must follow the host naming policy
2. there is no protection against re-using an old hostname that was registered in past but has been retired in the meanwhile, this complicates the tracking of hardware errors and failure analysis where access historical information is often required
3. We expect that VMs should address the majority of the use cases

Host aliases…

- the use of DNS aliases is supported and encouraged. They can be added by the responsible and the main user defined in LanDB

Special Cases/Requirements

Contact HW procurement team…

- https://cern.service-now.com/service-portal/report-ticket.do?name=request&fe=hw-procurement

Puppet-managed physical machines

Install and manage physical machines with Puppet and Foreman

- role: “physical machines”
  - request the role through via a Service-Now ticket.
- hardware allocated will become available in the **Incoming** hostgroup in Foreman.
  - the LanDB responsible will be able to edit host properties
  - installation process via PXE and for instance Anaconda and kickstart file

Partition table

In order to be installed, each physical machine must have a partition table defined.

- configured in the “Operating System” tab
- “Partition table” setting allows to select an existing template
  - Example: “Redhat default”
- to define your own configuration use disk text field
  - Partitioning information is provided using Anaconda/Kickstart partition commands

Partition table

Example of basic partitioning using volume groups and a single disk

```
clearpart    --drives sda --all
zerombr
part /boot   --size=1024 --ondisk sda
part pv.01   --size=1    --ondisk sda  --grow
volgroup vg1 pv.01
logvol /         --vgname=vg1 --size=10000  --name=root
logvol /var      --vgname=vg1 --size=10000  --name=var --grow
logvol /tmp      --vgname=vg1 --size=10000  --name=tmp
logvol swap  --vgname=vg1 --recommended --name=swap --fstype=swap
ignoredisk --only-use=sda
```

Partition table

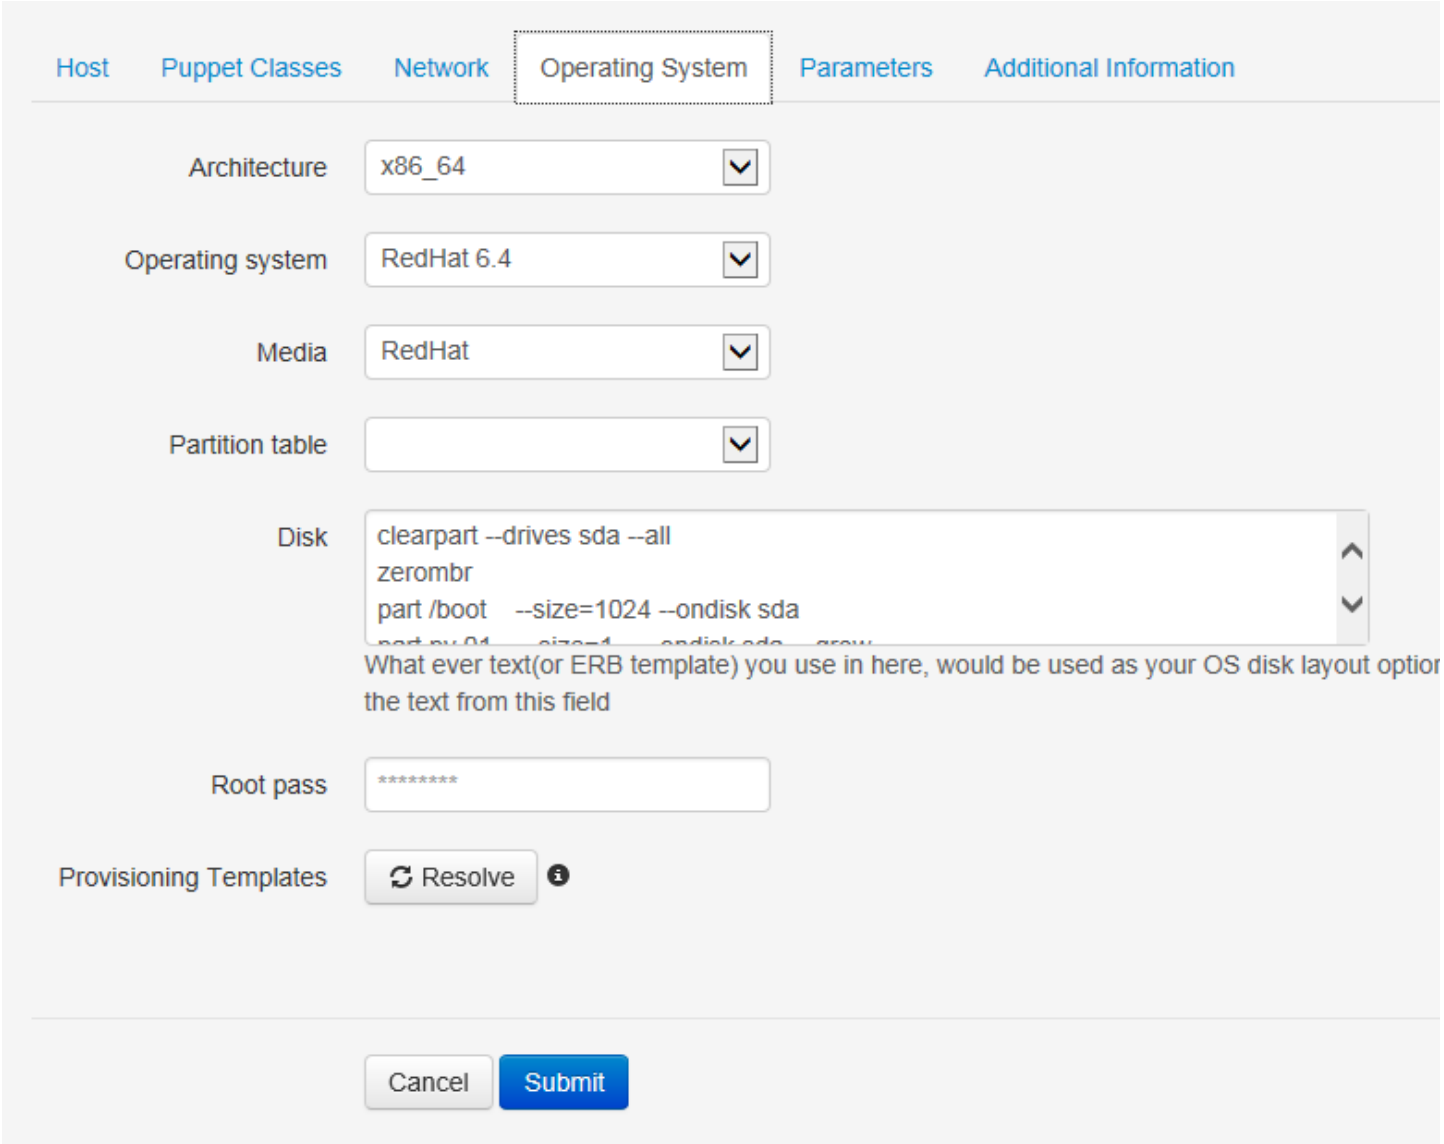

Kickstart provisioning templates

Provisioning Templates

Filter ... | Search | New Template | Build PXE Default

Name	Host group / Environment	Kind	Snippet	
A Kickstart Default	None / production and None / qa	provision		Delete
A Kickstart Default (TEST)	playground/alossent / qa	provision		Delete
CASTOR diskserver	castor/c2repack/diskserver/default / production, castor/c2cms/diskserver/t0cms / production, and castor/c2cms/diskserver/t0cms / qa	provision		Delete
CASTOR tapeserver		provision		Delete
diane_proxy_template		provision		Delete
diane_template		provision		Delete
EOS diskserver kickstart	eos/pps/storage / None, eos/public/storage / None, eos/atlas/storage / None, eos/alice/storage / None, eos/cms/storage / None, and eos/lhcb/storage / None	provision		Delete
epel			✔	Delete
Grubby Default		script		Delete
hamster_userdata		finish		Delete
hostcert			✔	Delete

Used to generate the kickstart file

- "A kickstart default" is provided and used by default and should be sufficient for most of the scenarios
- you should never modify or delete existing kickstart provising templates

Installing a physical machine

- Review the values set in the ‘Operating system’ tab:
  - architecture, OS and partition table template
- Prepare the machine for installation
  - From a terminal on aiadm.cern.ch

    ai-foreman-cli installhost --foreman-port 9443 myhostname
- Then reboot
  - ai-remote-power-control can be used to power on, off or power cycle Puppet-managed physical machines

    ai-remote-power-control cycle myhostname
  - BMC tab from Foreman
  - ‘Remote Control’
    - vendor-specific remote control interface
    - power operations
    - console access
- Problems and other operations on physical machines contact sysadmins team via an ITCM ticket

Edit Puppet-managed VMs in LandDB

- You cannot edit VMs in LandDB
  - for physical hardware, the owner of the machine can edit information in LanDB such as the main user or IP aliases
  - for virtual machines, this information is managed by OpenStack and thus the entry in LanDB for the virtual machines is locked from editing
  - If you need to modify entries, please open a ticket

    service desk portal

PuppetDB

- Collects data generated by Puppet and provides exported resources and inventory services in a Puppet Infrastructure

- PuppetDB's data can be queried with REST API, currently through a proxy using a valid CERN certificate

- Query structure

  https://judy.cern.ch:8081/v3/<ENDPOINT>?query=<QUERY STRING>

PuppetDB

1. querying with curl (facts of the vobox1 node)

curl -k --cert ./private/myCert.pem --key ./private/myCert.key -H 'Accept: application/json' https://judy.cern.ch:8081/v3/nodes/vobox1.cern.ch/facts

2. puppetDBQuery module
   - The Puppet module PuppetDBQuery provides functions to query PuppetDB directly from Puppet manifests.
   - Two functions available (query_nodes and query_facts)

Example: get the facts 'operatingsystem','osfamily‘ and 'architecture‘ of the vobox1 node

```
$variable= query_facts('hostname="vobox1"', ['operatingsystem','osfamily', 'architecture'])
```

More details about PuppetDB here

PuppetDB

3. ai-pdb

- script to query puppetdb using kerberos credentials
- installed in aiadm machines
- two subcommands “raw” and “hostgroup”
- more details man page ai-pdb

Example:

- get all the nodes where vobox is part of the name of the hostgroup

```
ai-pdb hostgroup vobox
ai-pdb raw /v3/nodes –query '["~",["fact","hostgroup"],"vobox"]'
```

The last VOC WG meeting…

We don't consider the VOC a special case anymore and from now on…

- There will be regular meetings and tutorials open to everyone…
  - every Thursday around 16h.
    - announced through ai-infrastructure-discuss mailing list
  - present new developments
  - you still can provide feedback
  - we invite to come to this regular meetings
    - there is one today ☺
- Future hands-on training will be organised
- The official documentation will cover everything you need to know about the configuration management process
- Questions and issues: configuration management support lines

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

??????????????????????

CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it