

# Cockpit Enhancements

## Enhancements in deployment flow

### Cache Setup Flow

This section described the enhancements in cockpit-gdeploy plugin to handle lvmcache configuration

1- In Gluster Deployment will have Check box based on Thinp check and will have one pop up.[Optional for User]

2- Below will be the fields which needs to be filled by User in pop up:

SSD: sdd

LV Size: XGB

-> If “Configure LV Cache” is unchecked then those two fields will be hidden.

The screenshot shows the 'Gluster Deployment' window with a progress bar at the top indicating the current step is 'Bricks' (step 3). Below the progress bar, the 'Raid Information' section is visible, showing 'RAID 6' selected for the Raid Type, '256' for Stripe Size(KB), and '12' for Data Disk Count. The 'Brick Configuration' section contains a table of bricks and a 'Configure LV Cache' section.

LV Name	Device Name	Size(GB)	Thinp	Mount Point
engine	sdb	100	<input checked="" type="checkbox"/>	/gluster_bricks/engine
data	sdb	500	<input type="checkbox"/>	/gluster_bricks/data
vmstore	sdb	500	<input checked="" type="checkbox"/>	/gluster_bricks/vmstore

Configure LV Cache [Add Bricks](#)

SSD:

LV Size(GB):

Cache Mode:

Caching mode is write-through by default and if cache is configured in other mode, please add input here.

Buttons: Cancel, <Back, Next>

Rest of the fields will be auto generated.

**3-** Then will make entry in gdeploy config file as:

```
[lv5]
```

```
action=setup-cache
```

```
ssd=sdd
```

```
vgname=gluster_vg_sdb or gluster_vg_vdb
```

```
poolname=gluster_thinpool_sdb or gluster_thinpool_vdb
```

```
cache_lv=lvcache
```

```
cache_lvsize=XGB
```

```
ignore_lv_errors=yes
```

```
cachemode=writethrough
```

**4-** Gdeploy will take care next.

## Enable brick/volume creation via Cockpit

Users should be able to provision additional hosts and additional bricks post the initial 3 node deployment.

Consider following cases:

1. Additional disks attached to the existing 3 nodes
2. Additional host added to a RHHI deployment
  - a. Host is added to RHV-M via UI, hosted-engine can be deployed on hosts while adding the host.

In both of these cases, the RHV-M UI has capability to create brick from a storage device. This is limited to creating a single LV on brick though. I.e - raw device -> PV -> VG -> Thinpool -> LV. In order to be able to achieve flexibility in terms of attaching LVM cache, enabling dedupe/ compression - adding this to RHV-M via blivet modules will take considerable time.

Two options:

1. allow disk provisioning on storage devices via Cockpit UI.
2. Provide UI in RHV-M to generate and call an Ansible role on the hosts that will do the brick provisioning.

If we were to go with option 1, Cockpit UI needs to support this.

## Complete Flow:

The screenshot displays the Red Hat Virtualization Host 4.1 (EL7.4) dashboard. The top navigation bar includes 'Virtualization' and 'Dashboard' tabs. A sidebar on the left contains 'Dashboard', 'Hosted Engine', and 'Virtual Machines' sections. The main content area shows the status of the hosted engine and a list of hosts in the cluster.

**Status of this host (rhsqa-grafton1.lab.eng.blr.redhat.com)**

Hosted Engine is up! Hosted Engine is running on rhsqa-grafton1.lab.eng.blr.redhat.com

Actions: Put this host into local maintenance | Remove this host from maintenance | Put this cluster into global maintenance

Hosts in this cluster		Manage Cluster
<b>rhsqa-grafton1.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: up	<b>rhsqa-grafton2.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false
<b>rhsqa-grafton3.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: down	<b>VM Status</b> State: down

## Gluster Management:

### Gluster Management

---

#### HOSTS

Name	Peer Status
Host 1	Peer in cluster (connected)
Host 2	Peer in cluster (connected)
Host 3	Peer in cluster (connected)

---

#### VOLUMES

[Expand Cluster](#)

Name	Cluster	Volume Type	Bricks								
Volume 1	Default	Replicate	3  0								
		<table><thead><tr><th>Name</th><th>Brick Directory</th></tr></thead><tbody><tr><td>brick1</td><td>/gluster-bricks/brick1/brick1</td></tr><tr><td>brick2</td><td>/gluster-bricks/brick2/brick2</td></tr><tr><td>brick3</td><td>/gluster-bricks/brick3/brick3</td></tr></tbody></table>	Name	Brick Directory	brick1	/gluster-bricks/brick1/brick1	brick2	/gluster-bricks/brick2/brick2	brick3	/gluster-bricks/brick3/brick3	<a href="#">Expand Volume</a>
Name	Brick Directory										
brick1	/gluster-bricks/brick1/brick1										
brick2	/gluster-bricks/brick2/brick2										
brick3	/gluster-bricks/brick3/brick3										
Volume 2	Default	Replicate	3  0								
Volume 3	Default	Replicate	2  1								

[Create Volume](#)

-> This page will have Hosts, Bricks and Volumes informations.

-> From here we can Expand Cluster, Expand Volume and Create Volume.

## Expand Cluster:

Gluster Management ✕

Hosts Packages Volumes Bricks Review

1 2 3 4 5

Host1

Host2

Host3

ⓘ Make sure, passwordless ssh is configured for the host.

-> Here we can expand our cluster by adding new Hosts.

The image shows a window titled "Gluster Management" with a close button (X) in the top right corner. At the top, there is a progress bar with five steps: Hosts (1), Packages (2), Volumes (3, highlighted in blue), Bricks (4), and Review (5). Below the progress bar, there are four input fields: "Name" (empty), "Volume Type" (a dropdown menu showing "Replicate"), "Arbiter" (a checkbox that is unchecked), and "Brick Dirs" (empty). Below these fields is a blue button labeled "Add Volume". At the bottom right of the window, there are three buttons: "Cancel", "< Back", and "Next >" (highlighted in blue).

**-> Initially will have one volume, for more we can add more volumes by clicking on "Add Volume".**

Gluster Management ✕

Hosts Packages Volumes Bricks Review

① — ② — ③ — ④ — ⑤


**Raid Information** ⓘ

Raid Type: RAID 6 ▾

Stripe Size(KB):

Data Disk Count:

**Brick Configuration**

LV Name	Device Name	Size(GB)	Thinp	Mount Point
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>	<input type="text"/> 

[⊕ Add Bricks](#)

Cancel < Back Next >

-> We can configure more bricks by using “Add bricks”

RED HAT VIRTUALIZATION HOST 4.1 (EL7.4)

Virtualization Dashboard

Hosted Engine is up! Hosted Engine is running on rhsqa-grafton1.lab.eng.blr.redhat.com

Status of this host (rhsqa-grafton1.lab.eng.blr.redhat.com)

rhsqa-grafton1.lab.eng.blr.redhat.com

Put this host into local maintenance | Remove this host from maintenance | Put this cluster into global maintenance

Hosts in this cluster Manage Cluster

<b>rhsqa-grafton1.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: up	<b>rhsqa-grafton2.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: down
<b>rhsqa-grafton3.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: down	<b>rhsqa-grafton4.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: down
<b>rhsqa-grafton5.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: down	<b>rhsqa-grafton6.lab.eng.blr.redhat.com</b> Agent stopped: false Local Maintenance: false	<b>VM Status</b> State: down

-> This will be the final screen after expand cluster.



## Expand Volume:

The image shows a 'Gluster Management' dialog box with a progress indicator at the top. The progress indicator consists of four steps: 'Hosts' (1), 'Volumes' (2), 'Bricks' (3), and 'Review' (4). The 'Hosts' step is currently selected and highlighted with a blue circle. Below the progress indicator, there are three rows of host selection. Each row has a label ('Host1', 'Host2', 'Host3') followed by a dropdown menu containing the text '-- Select Host --'. At the bottom right of the dialog box, there are three buttons: 'Cancel', '< Back', and 'Next >'. The 'Next >' button is highlighted in blue.

-> For expand volume, will have to choose hosts present in cluster.

Gluster Management

Hosts Volumes Bricks Review

1 2 3 4

Name: -- Select Volume -- Volume Type: Replicate Arbitrator:  Brick Dirs:

Cancel < Back Next >

-> Here will have to choose volume and after that all related value will be populated which are readonly.

Gluster Management X

Hosts      Volumes      **Bricks**      Review

① ——— ② ——— ③ ——— ④

---

**Raid Information** ⓘ

Raid Type:

Stripe Size(KB):

Data Disk Count:

**Brick Configuration**

LV Name	Device Name	Size(GB)	Thinp	Mount Point
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>	<input type="text"/>

## Create Volume:

Gluster Management ✕

Hosts      Volumes      Bricks      Review

1 — 2 — 3 — 4

Host1  ▾

Host2  ▾


Host3  ▾

Gluster Management ✕

Hosts      Volumes      Bricks      Review

① ——— ② ——— ③ ——— ④

---

Name       Volume Type Replicate ▾      Arbiter       Brick Dirs  

[⊕ Add Volume](#)

---



**Raid Information** ⓘ

Raid Type:  ▾

Stripe Size(KB):

Data Disk Count:

**Brick Configuration**

LV Name	Device Name	Size(GB)	Thinp	Mount Point
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>	<input type="text"/> 

[+ Add Bricks](#)

## Asymmetric brick configuration:

The screenshot shows the 'Bricks' configuration step in the Gluster Deployment wizard. The wizard has five steps: Hosts (1), Packages (2), Volumes (3), Bricks (4), and Review (6). The 'Bricks' step is currently active.

**Raid Information**

- Raid Type: RAID 6
- Stripe Size(KB): 256
- Data Disk Count: 12

**Brick Configuration**

Hosts: Host1, Host2 (selected), Host3

LV Name	Device Name	Size(GB)	Thinp	Mount Point
engine	sdb	100	<input checked="" type="checkbox"/>	/gluster_bricks/engine
data	sdb	500	<input type="checkbox"/>	/gluster_bricks/data
vmstore	sdb	500	<input checked="" type="checkbox"/>	/gluster_bricks/vmstore

Configure LV Cache

- SSD: sdc
- LV Size(GB): 5
- Cache Mode: writethrough

By default cachemode is writethrough and if cache is configured in writeback mode then please add input here.

Buttons: Cancel, <Back, Next>

-> In brick configuration will have “Tab” for Hosts and Host wise you can change Device.

-> Default all the fields will be read only except “Device Name”

-> If the volume is “Arbiter” then “Size” field will be writable.

- Changing Device in Host3 and also volume is “Arbiter”:

Gluster Deployment X

Hosts Packages Volumes Bricks Review

① — ② — ③ — ④ — ⑥

**Raid Information** ⓘ

Raid Type: RAID 6  
Stripe Size(KB): 256  
Data Disk Count: 12

**Brick Configuration**

Host1 Host2 **Host3**

LV Name	Device Name	Size(GB)	Thinp	Mount Point
engine	sdc	100	<input checked="" type="checkbox"/>	/gluster_bricks/engine
data	sdc	500	<input type="checkbox"/>	/gluster_bricks/data
vmstore	sdc	500	<input checked="" type="checkbox"/>	/gluster_bricks/vmstore

[Add Bricks](#)

Configure LV Cache

SSD: sdc  
LV Size(GB): 5  
Cache Mode ⓘ: writethrough

By default cachemode is writethrough and if cache is configured in writeback mode then please add input here.

Cancel <Back Next>

-> Here Host3 device is “sdc”.