

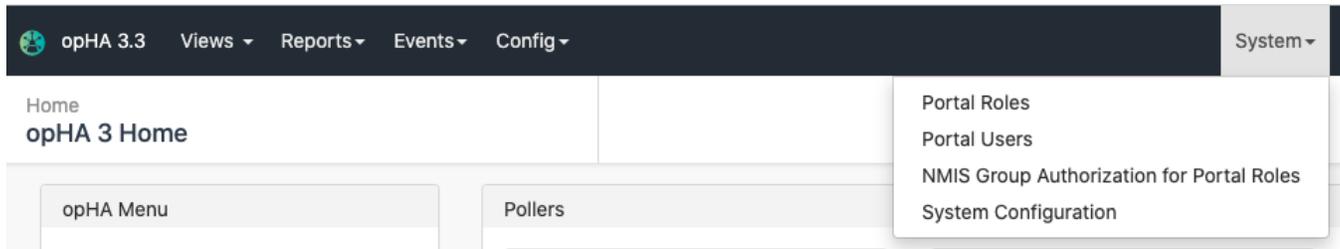
opHA 3 Redundant Node Polling and Centralised Node Management

- Centralised Node Management
 - Technical details
 - Managing remote nodes
- Redundant Node Polling
 - Architecture Schema
 - Discoveries
 - What happens when we discover a poller?
 - What happens when we discover a mirror?
 - Deletion
 - What happens when we remove a poller?
 - What happens when we discover a mirror?
 - Pulling Data
 - What happens when you pull from a poller?
 - What happens when you pull from a mirror?
 - Data Verification
 - opHA Configuration

Centralised Node Management

This new feature will allow us to manage all remote nodes from the main server.

The only thing you need to do is to access to the new **System Configuration admin** menu:

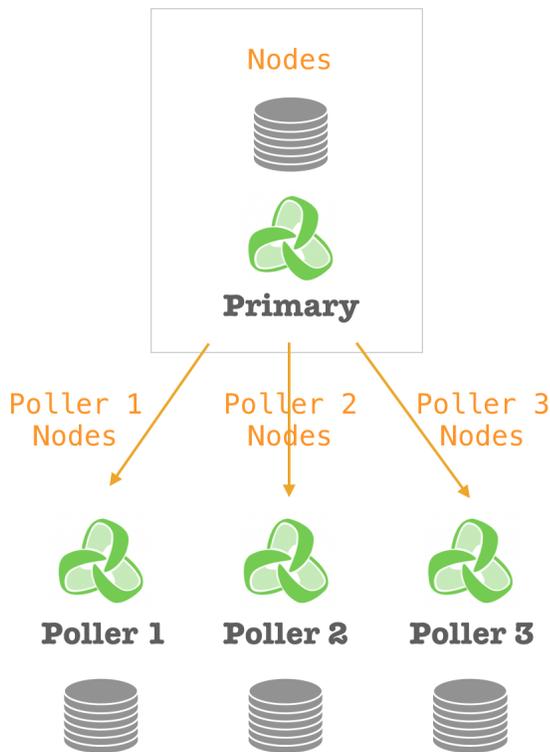


<http://server/en/omk/admin>

Home / Nodes
Nodes

System Management								
Nodes								
<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>								
<input type="text" value="Search term"/> <input type="text" value="Name"/> <input type="button" value="Go"/> <input type="button" value="X"/>								
<input type="checkbox"/>	Name ▲	Group	Location	Business Service	Host	Addresses	Activated	Cluster
<input type="checkbox"/>	ASGARD	DataCentre	Gold Coast	Web Page	192.168.88.254			fulla
<input type="checkbox"/>	Adtran-TA5000	SNMPSIM	Amsterdam		10.152.0.39			fulla
<input type="checkbox"/>	Angmar	123	Amsterdam		138.68.11.77			fulla
<input type="checkbox"/>	AroValley	Branches	Aro		210.54.38.168			fulla
<input type="checkbox"/>	Axtell	SNMPSIM	Amsterdam		10.152.0.39			fulla
<input type="checkbox"/>	Beattie	SNMPSIM	Amsterdam		10.152.0.39			fulla
<input type="checkbox"/>	C880	SNMPSIM	Amsterdam		10.152.0.39			fulla
<input type="checkbox"/>	Ciena-Waveserver	NMIS9			Ciena-Waveserver			Local
<input type="checkbox"/>	Ciena-Waveserver	SNMPSIM	Amsterdam		10.152.0.39			fulla
<input type="checkbox"/>	Ciena-Waveserver	dodgy&group	Cloud		10.152.0.39			poller-nine
<input type="checkbox"/>	CiscoSR4451	SNMPSIM	Amsterdam		10.152.0.39			fulla
<input type="checkbox"/>	HY-COM008-0004-RTR-MT-1-EAST-LONDON-AIRPORT	123	Aro		100.64.1.37			fulla

Technical details



- New administration menu to manage nodes. The GUI will work differently based on the server role:
 - Standalone server: Can manage local nodes.
 - Poller: Just view nodes.
 - Primary server: Can manage local and remote nodes. This means, the Primary server can replicate the information to the remote pollers.

Managing remote nodes

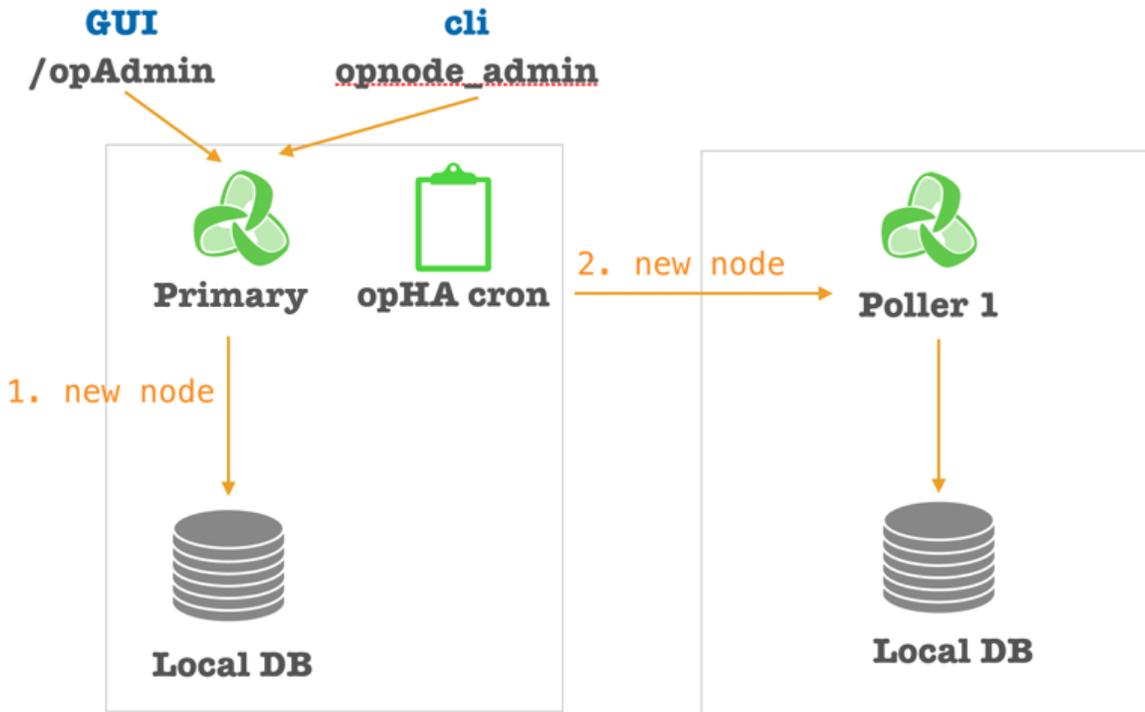
There are several tools and processes to work with remote nodes:

Tool	Access	Working mode	Notes
Admin GUI	<ul style="list-style-type: none"> • http://server/en/omk/admin/nodes 	Online	Will sync the nodes online
CLI Tool	<code>/usr/local/omk/bin/opnode_admin</code>	Background	See the documentation here
CLI Tool	<code>/usr/local/omk/bin/oph-cli.pl act=sync-all-nodes</code>	Online cron job - 5 minutes	Will review and sync all the nodes
CLI Tool	<code>/usr/local/omk/bin/oph-cli.pl act=sync-node</code>	Online	Will sync just one node
CLI Tool	<code>/usr/local/omk/bin/oph-cli.pl act=sync-processed-nodes</code>	Background cron job - 5 minutes	Will sync the nodes processed by opnode_admin
CLI Tool	<code>/usr/local/omk/bin/oph-cli.pl act=resync_nodes peer=server-name</code>	Online	Will remove the nodes from the poller in the Primary will pull the nodes from the poller

Why different processes?

- **sync-all-nodes:** Is a more robust process. This will check what the remote nodes have and will review what needs to be updated based on the local database, with the Primary as a source of truth. That means, if the node does not exist in the poller, it will create the node in the poller. If the node does not exist in the Primary, it will be removed from the poller. If the information is not the same (based in the last updated field), it will be updated in the poller. This is more robust, but less efficient. It has to do a remote call to get all nodes from the pollers, and check all of the nodes one by one. **Important notes:**
 - If the Primary cannot get the nodes from the poller, the synchronisation will not be done.
 - If the nodes list in the Primary for one poller is 0, the remote nodes won't be removed. As it could mean that the nodes were not correctly synchronised in the beginning, or there was an issue trying to get the Primary nodes list.

- **sync-processed-nodes:** This process will update all nodes on remote pollers based on the information updated by opnode_admin. This is a less robust process, as all the changes done outside opnode_admin won't be taken into account, but it is more effective, as it will sync just the data needed.
- **opnode_admin:** Will update the information to be processed in background - So, it is a non blocking operation.
- **resync_nodes:** In case it is necessary to pull the nodes from the poller for some reason (Discovery failed to bring initial nodes, ex).



Redundant Node Polling

With this new feature, opHA can set up a mirror server for a poller. The mirror server will be polling the same devices than the poller.

opHA Menu

- Peers**
Discover and manage peers
- Configuration**
Edit remote configuration files
- Log**
Peers last activity
- Verify**
Data verification

Pollers

fulla	
✓	
Last pull	Mon May 24 22:01:52 2021
Pull Status	Success
Nodes	81
DB status	✓
opchartsd	Error
opevents	Ok
opconfigd	Error
nmis9d	Ok
omkd	Error

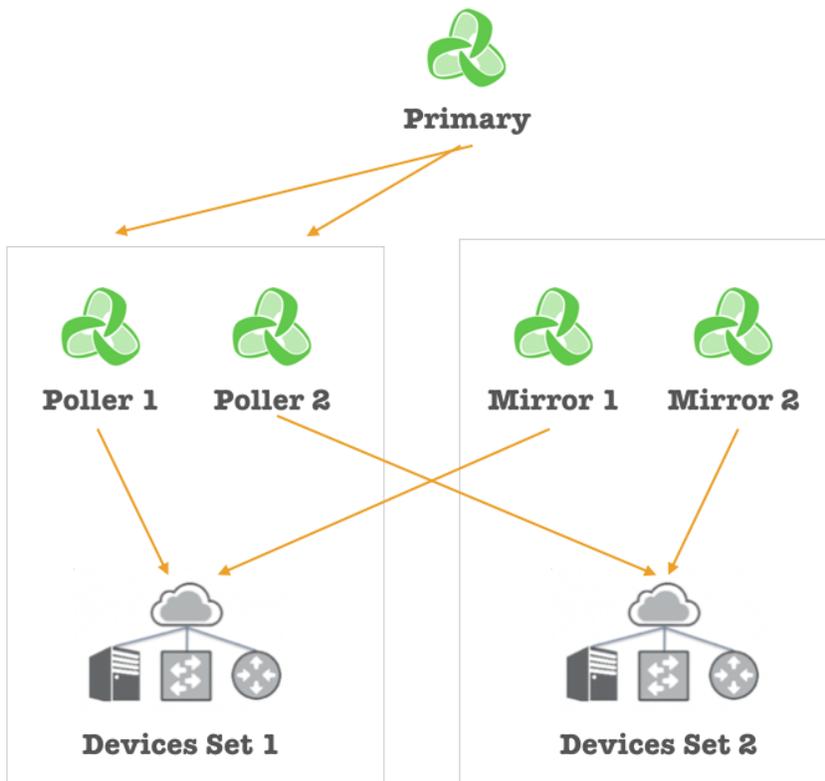
uburnto - (fulla Mirror)	
✓	
Last pull	Tue May 25 15:15:51 2021
Pull Status	Success
Nodes	
DB status	✓
opchartsd	Ok
nmis9d	Ok
omkd	Ok

Primary

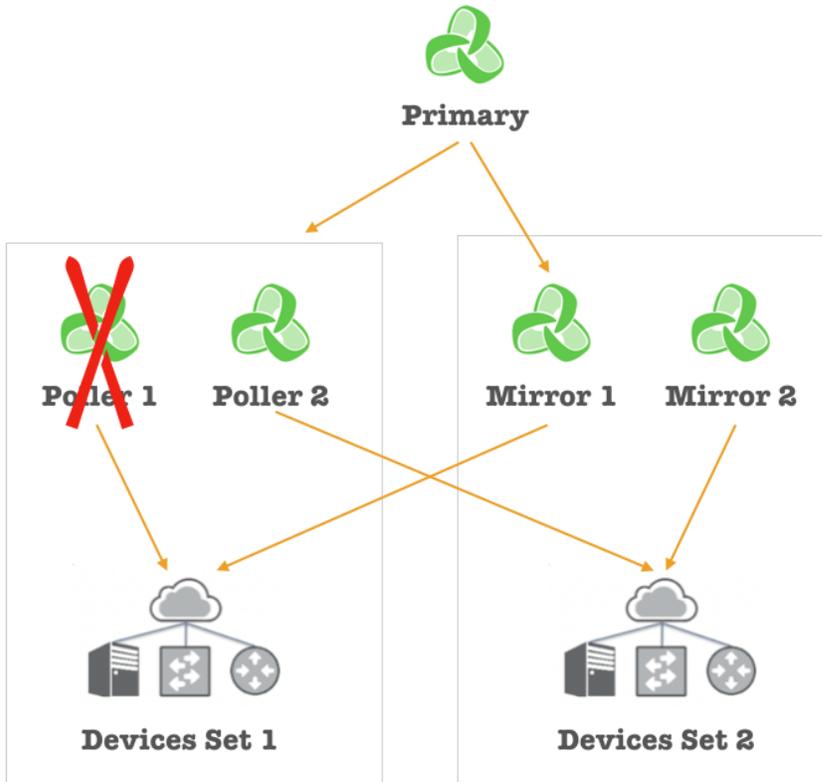
Local	
✓	
Last pull	-

[Here](#) you can find a more detailed information with the roles and functionality of each one.

Architecture Schema



When poller 1 is offline, the Primary will collect all the data from the Mirror 1 automatically.



It is a requirement that the devices collected from the poller are the same than the mirror.

Discoveries

What happens when we discover a poller?

- The first time we discover a poller, all the nodes are copied into the primary. But once they have been copied, in the synchronisations the node configuration will be sent from the primary to the poller.
- The server will be set up as a poller, so we are not gonna be able to edit the nodes from the poller.

What happens when we discover a mirror?

- The first time we discover a mirror, it needs to be empty (Except for localhost). This is because, all nodes needs to be copied from the poller that is being mirrored.
- The nodes will be sent to the mirror, the mirror will start polling those devices, at the same time as the poller.
- But the data won't be synchronised, unless the poller is disabled/down.
- The server will be set up as a mirror, so we are not going to be able to edit the nodes from the mirror.
- Every time we make a change into the poller nodes from the primary, they will be replicated to the mirror.

Deletion

What happens when we remove a poller?

We should update the role before removing the poller. If not, we are not going to be able to manage the nodes. We can do this manually by setting `opha_role` in `conf/opCommon.json` and restarting the server.

The mirror will become the poller.

What happens when we discover a mirror?

We can update the role before removing the poller. If not, we are not going to be able to manage the nodes. We can do this manually by setting `opha_role` in `conf/opCommon.json` and restarting the server.

Pulling Data

What happens when you pull from a poller?

If the poller is active and responsive, we will bring the data from the poller, except for the nodes. This is made in the **Sync all Nodes** operation.

What happens when you pull from a mirror?

If the poller is active, we will only bring the registry and the selftest data.

Data Verification

We can use the following command to review the data and where it is coming from:

```
/usr/local/omk/bin/ophi-cli.pl act=data_verify debug=9
```

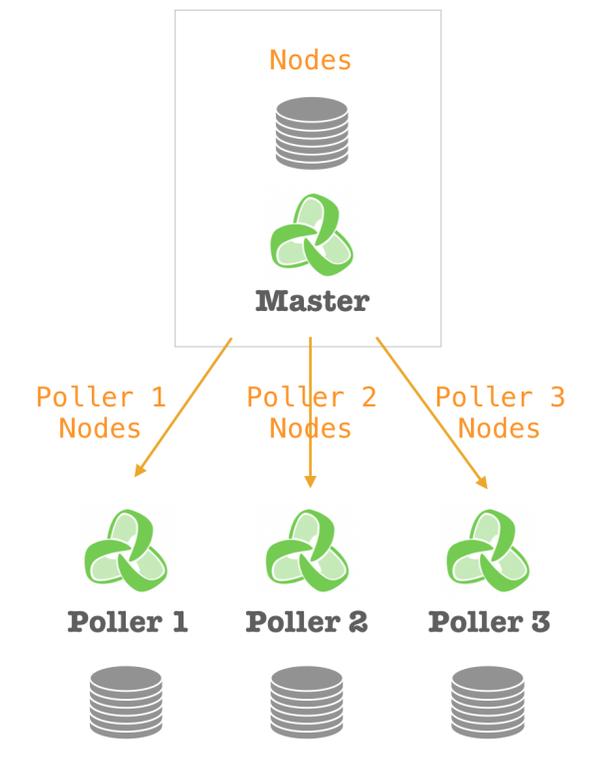
As an output example:

```
ACTIVE ENABLED STATUS NODES INVENTORY EVENTS LATEST_DATA STATUS
=====
Poller: poller-nine 0d28dcf0-8fe2-49d9-a26f-6ccf3f2875c0
      1      ok      [ 519 ]      [ 20442 ]      [ 694 ]      [ 3956 ]      [ 2059 ]

Poller: fulla a515c63a-0d11-4bcb-a402-39a4600aebb9
      1      ok      [ 99 ]      [ 2827 ]      [ 2827 ]      [ ]      [ ]

Poller: uburnto 8bf31ddc-8f32-4fdf-8173-8ce989e2e210
      1      ok      [ 7 ]      [ 230 ]      [ 10 ]      [ 40 ]      [ 101 ]
[Tue May 11 19:40:09 2021] [info] Returning: se7en
Mirror: se7en 1487b8fb-f1f9-41e9-995b-29e925724ff3
      0      1      ok      [ ]      [ ]      [ ]      [ ]
```

Also, a new GUI menu has been introduced to show the same data:



Data verification GUI:

Peer Data

Server Name	Cluster Id	Active	Enabled	Nodes	Inventory	Events	Latest Data	Status	Role	Mirror	DB Status
fulla	a515c63a-0d11-4bcb-a402-39a4600aebb9	1	1	81	3017	2792	587	775	Poller		ok
uburnto	8bf31ddc-8f32-4fdf-8173-8ce989e2e210	0	1						Mirror	a515c63a-0d11-4bcb-a402-39a4600aebb9	ok

<< < 1 > >>

Show 5

opHA Configuration

Hostname: volla
URL:

Both properties (opha_hostname and opha_url_base) need to be set to generate the registry

How to know if the data is correct?

- If the role is a mirror, needs to have a mirror associated.
- The mirror associated should exist.
- Mirror and poller should not be "Active" at the same time.
- "Active" is set internally in the system and is a flag for the poller and the mirror. But does not need to be set in poller with no mirror.
- For mirror and poller: The active one should show the number of data, none for the other.
- The mirror doesn't have to have Nodes in the report. These nodes are counted from the primary database. But, the remote should have the same nodes that the poller has.

opHA Configuration

We will see a warning when just one of the properties, opha_hostname and opha_url_base, are set. Both can be unset, or with values, to have a proper configuration.