

opCharts Inventory Database Performance

This feature is specifically for scaling inventory display in massive environments with large inventory counts. In environments with large inventory counts the inventory page can get bogged down when trying to join records together and calculating pagination.



This requires opCharts-4.5.2 and greater NMIS-9.4.2, with opHA you will need all servers in the cluster to run NMIS-9.4.2 and greater installed.

Enable the Feature

Update the OMK configuration file, /usr/local/omk/conf/opCommon.json

opCommon.json

```
"opcharts_inventory_use_cached_node_records" : 1
```

Restart the OMK daemon

```
sudo systemctl restart omkd
```

Update the Inventory

The update to NMIS9 will update the inventory automatically, if you want to wait for that to run on all nodes, to make the update now, use the cli command to back-fill the inventory records on the primary and pollers.

```
/usr/local/omk/bin/opcharts-cli.pl act=enhance-inventory
```

This will add the node name and node group to the inventory records, note if you change the nodes group or rename it could cause unexpected results in the inventory screen. In this case you can run the enhance inventory cli command again to re-patch the records.

The following fields are included:

- node_name
- configuration.group

Functionality

If the inventory page is trying to display over 100,000 records the gui will skip counting the records and the Display Grid will swap into a infinite mode which will paginate through the records until it hit the end. The count will show (pageSize * pageNum) of Many.

Skipping this total count results in better query performance and skipping MongoDB having to count all the records.

Troubleshooting

Clear the cache

If your inventory is incomplete, refresh the inventory cache.

```
/usr/local/omk/bin/opcharts-cli.pl act=cache-resources force=1
```