

# Configuration Options for Server Performance Tuning

There are lots of factors that determine the system health of a server. The hardware capabilities - CPU, memory or disk - is an important one, but also the server load - number of devices (Nodes to be polled, updated, audited, synchronised), number of products (NMIS, OAE, opCharts, opHA - each running different processes), number of concurrent users.

We all want the best performance for a server, and to optimise physical resources, our configuration has to be fine-grained adjusted. In this guide you will find recommended parameters, that may not suit in all cases, as a server performance will depend on a lot of factors.

- [Related Articles](#)
- [Opmantek Applications](#)
  - [Before Start](#)
  - [Configuration items](#)
  - [Server examples](#)
    - [Stressed system POLLER-NINE](#)
    - [Healthy system MASTER-NINE](#)
    - [Stressed system CUSTOMER SERVER UZH](#)

## Related Articles

- [Scaling NMIS Polling](#)
- [Scaling NMIS polling - how NMIS handles long running processes](#)
- [NMIS 8 - Configuration Options for Server Performance Tuning](#)
- [NMIS 9 - Configuration Options for Server Performance Tuning](#)
- [opCharts 3 Performance Tuning](#)

## Opmantek Applications

This article configurations are related to Opmantek products. opCharts, opEvents, opConfig, opHA, opReports, ... all use the omkd daemon which servers the frontend requests. Also, opEvents, [opCharts](#) and opConfig have their own daemons.

### Before Start

The first thing to do will be get the information of our system:

- **System Information:** NMIS and OMK [support tool](#) will give us all the information needed.
- **Monitor services:** NMIS can [monitor the involved processes](#) - apache2, nmis9d, omkd and mongod - and provide useful information about CPU and memory - among others.

### Configuration items

In low memory environments lowering the number of omkd workers provides the biggest improvement in stability, even more than tuning mongod.conf does. The default value is 10, but in an environment with low users concurrency it can be decreased to 3-5.

```
omkd_workers
```

Setting also omkd\_max\_requests, will help to have the threads restart gracefully before they get too big.

```
omkd_max_requests
```

Process size safety limiter: if a max is configured and it's >= 256 mb and we're on linux, then run a process size check every 15 s and gracefully shut down the worker if over size.

```
omkd_max_memory
```

Process maximum number of concurrent connections, defaults to 1000:

```
omkd_max_clients
```

The performance logs are really useful for debugging purposes, but they also can affect performance. So, it is recommended to turn them off when they are not necessary:

```
omkd_performance_logs => false
```

## MongoDB memory usage

MongoDB, in its default configuration, will use will use the larger of either 256 MB or  $\frac{1}{2}$  of (ram – 1 GB) for its cache size.

MongoDB cache size can be changed by adding the **cacheSizeGB** argument to the `/etc/mongod.conf` configuration file, as shown below.

```
storage:
  dbPath: /var/lib/mongodb
  journal:
    enabled: true
  wiredTiger:
    engineConfig:
      cacheSizeGB: 1
```

Here is an interesting information regarding how MongoDB reserves memory for internal cache and WiredTiger, the underneath technology. Also some adjustment that can be done: <https://dba.stackexchange.com/questions/148395/mongodb-using-too-much-memory>

## Server examples

Two servers are compared in this section.

- Primary only have one node, but more than 400 poller nodes. opHA process is what will require more CPU and memory usage.
- Poller have more more than 500 nodes. nmis process will require more CPU and memory, for polling the information for all the nodes.

### Stressed system

POLLER-NINE


System information:

Name	Value	Notes
nmisd_max_workers	10	(nmis9 only)
omkd_workers	4	
omkd_max_requests	500	
Nodes	406	
Active Nodes	507	
OS	Ubuntu 18.04.3 LTS	
role	poller	

This is how the server memory graphs looks in a stressed system - We will focus on the memory as this is where the bottleneck is:


## Real Memory Usage

Start	04-Nov-2019 09:56:18	Node	poller-nine	Type	Host_Memory	Submit
End	06-Dec-2019 09:56:18	Interface		<a href="#">NMIS IP Response Health Stats Export Adv. Export</a>		



Clickable graphs: Left -> Back; Right -> Forward; Top Middle -> Zoom In; Bottom Middle-> Zoom Out, in time

Start	04-Nov-2019 09:57:21	Node	poller-nine	Type	hrcachemem	Submit
End	06-Dec-2019 09:57:21	Interface		<a href="#">NMIS Response IP Health Stats Export Adv. Export</a>		




Clickable graphs: Left -> Back; Right -> Forward; Top Middle -> Zoom In; Bottom Middle-> Zoom Out, in time


NMIS process remains stable, is not using more than 120 mb, and the process was stopped - probably killed for the system due to high memory usage:

TODO


How to check this



The OMK process has more fluctuations and higher memory usage - peaks up to 800 mb. The memory trend is to raise:



And mongod keeps using a lot of memory - 3GB, as configured - but it is stable:



Check processes once nmis9d is restarted again:

top

```
top - 11:12:30 up 92 days, 11:45, 1 user, load average: 8.22, 2.85, 1.07
Tasks: 140 total, 13 running, 85 sleeping, 0 stopped, 1 zombie
%Cpu(s): 84.7 us, 12.8 sy, 0.0 ni, 0.2 id, 1.3 wa, 0.0 hi, 1.0 si, 0.0 st
KiB Mem : 7652360 total, 146312 free, 2934448 used, 4571600 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 4448044 avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
20219	mongodb	20	0	2469712	1.410g	31080	S	23.3	19.3	15:14.93	mongod
1432	root	20	0	353932	124036	9160	R	16.9	1.6	0:10.56	nmisd worker co
1439	root	20	0	353932	124020	9160	R	15.6	1.6	0:09.61	nmisd worker co
1441	root	20	0	353932	124036	9160	R	15.6	1.6	0:09.61	nmisd worker co
1434	root	20	0	353932	124020	9160	R	15.3	1.6	0:09.74	nmisd worker co
1426	root	20	0	357648	121804	7536	R	15.0	1.6	0:07.88	nmisd fping
1446	root	20	0	353932	124036	9160	R	15.0	1.6	0:09.70	nmisd worker co
1448	root	20	0	353932	124036	9160	R	15.0	1.6	0:09.62	nmisd worker co
1430	root	20	0	353980	124020	9160	R	14.0	1.6	0:11.33	nmisd worker co
1428	root	20	0	354048	124216	9228	R	13.3	1.6	0:18.51	nmisd worker co
1443	root	20	0	353920	124036	9160	R	13.3	1.6	0:10.33	nmisd worker co
1437	root	20	0	353932	124036	9160	R	13.0	1.6	0:09.83	nmisd worker co
1425	root	20	0	355396	121296	5452	R	5.3	1.6	0:10.42	nmisd scheduler
125	root	0	-20	0	0	0	I	3.3	0.0	187:58.56	kworker/0:1H
1520	root	20	0	0	0	0	Z	1.0	0.0	0:00.03	fping
41	root	20	0	0	0	0	S	0.7	0.0	408:36.96	kswapd0
124	root	0	-20	0	0	0	I	0.7	0.0	68:30.16	kworker/1:1H



Healthy system

MASTER-NINE

System information:

Name	Value
nmisd_max_workers	5
omkd_workers	10
omkd_max_requests	undef
Nodes	2
Poller Nodes	536
OS	Ubuntu 18.04.3 LTS
role	master

This is how the server memory graphs looks in a normal system:



Daemons graphs:

## Service Memory

Start 04-Nov-2019 11:20:27

Node master-nine


Type service-mem

Submit

End 06-Dec-2019 11:20:27

Service nmisd

[NMIS Health](#) [Response](#) [IP](#) [Stats](#) [Export](#) [Adv.Export](#)



Clickable graphs: Left -> Back; Right -> Forward; Top Middle -> Zoom In; Bottom Middle-> Zoom Out, in time

omk:

## Service Memory

Start 04-Nov-2019 11:20:49

Node master-nine


Type service-mem

Submit

End 06-Dec-2019 11:20:49


Service opmantek

[NMIS Response](#) [Health](#) [IP](#) [Stats](#) [Export](#) [Adv.Export](#)



Clickable graphs: Left -> Back; Right -> Forward; Top Middle -> Zoom In; Bottom Middle-> Zoom Out, in time

mongo:



Stressed system

CUSTOMER SERVER UZH

System information:


Name	Value
nmisd_max_workers	50
nmisd_scheduler_cycle	30
nmisd_worker_cycle	10
nmisd_worker_max_cycles	10

nmisd is crashing with no error messages.

Some server info:

- CentOS 7
- 463 Nodes
- Poller server
- High IO Wait





- increased open files to 100'000