OPMANTEK

Understanding the NMIS KPI interface

E-book by

Daniel Carter
Marketing Coordinator

Table of Contents

What is a KPI and why is it relevant it for network monitoring?

Key Performance Indicators (KPIs) were introduced into NMIS to provide insight as to why the health of a node was getting better or worse. As discussed in the article on **NMIS Metrics, Reachability, Availability and Health**, NMIS is tracking the health of a node and providing a single number which indicates what the health of a node is, this is called the Health Metric. To make up the Health Metric, NMIS is tracking many aspects of a node's health including:

- Reachability – Node availability or pingability
- Availability – Interface availability
- Response time
- CPU Utilisation
- Memory Utilisation
- Interface Utilisation
- Disk Utilisation
- Swap Utilisation

NOTE: N*ot all nodes have disk and swap, so for some nodes these values are blank, e.g. a Cisco Router will have no value for disk and swap KPI's.*

NMIS has a history of being a Network Management System, the generation of the Metrics and KPI's is something that makes NMIS more than a Network Monitoring System and helps IT professionals by providing better information about their environment to help with their decisions. By giving users more information about devices, troubleshooting or improving the health of devices is much easier, this can even be extended further by adding opTrend, which gives metrics using machine learning to build time-based device expectations.

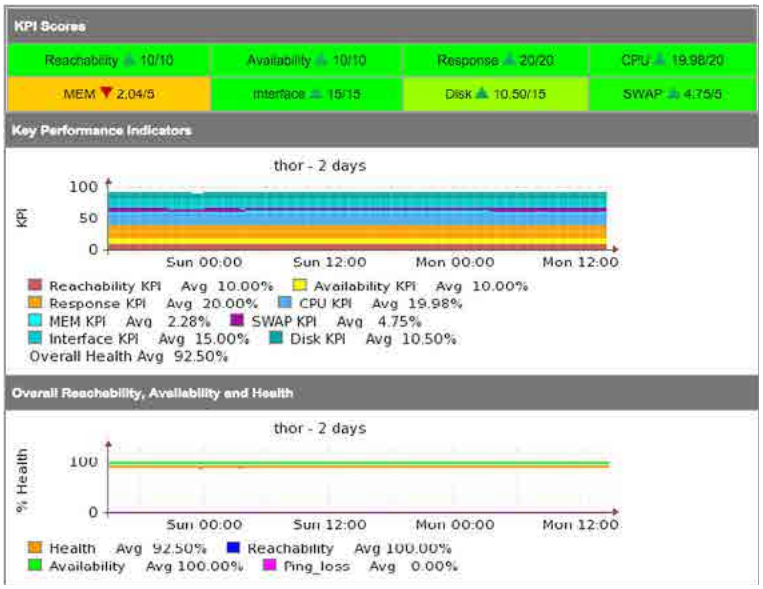

As of NMIS 8.5G, we started storing the individual KPI scores so that it was possible to see the health metric break down over time. This is now shown at the top of a node view panel in NMIS8 and looks like the image above.

KPI Scores

You can think of the KPI Scores like a report card, the student (node) has received 10/10 for English (reachability), 10/10 for Maths (availability) and so on. The KPI Scores in the screenshot above come from the polled data and are scored out of the weighted value, this weighted value is a percentage, so in the configuration file, it is 0.1 which means it is 10% or a maximum possible KPI score of 10/10. The table below shows the configuration value and the resulting KPI Score value.

KPI ITEM	CONFIGURATION ITEM	CONFIGURED WEIGHTING	MAXIMUM KPI SCORE
Reachability	weight_reachability	0.1	10 (10%)
Availability	weight_availability	0.1	10 (10%)
Response	weight_response	0.2	20 (20%)
CPU	weight_cpu	0.2	20 (20%)
Memory	weight_mem	0.1	10 (10%)
Interface	weight_int	0.3	30 (30%)

Because they are not present in all node types, there are two additional KPI values which overload onto the Memory and Interface KPI values these are, Swap and Disk, these split the weighting of each into half and track that separately, e.g. Interface KPI by default is 30%, so when the Disk KPI is present the Interface KPI gets a value of 15% and the Disk KPI gets a value of 15%. So the table would like this when all 8 KPI's are present, as they are for Linux Servers.

ITEM	CONFIGURATION ITEM	CONFIGURED WEIGHTING	MAXIMUM KPI SCORE
Reachability	weight_reachability	0.1	10 (10%)
Availability	weight_availability	0.1	10 (10%)
Response	weight_response	0.2	20 (20%)
CPU	weight_cpu	0.2	20 (20%)
Memory	weight_mem	0.1 x 50%	5 (5%)
Swap	weight_mem	0.1 x 50%	5 (5%)
Interface	weight_int	0.3 x 50%	15 (15%)
Disk	weight_int	0.3 x 50%	15 (15%)

The result is that all the maximum KPI Score for a node will be 100 or 100%.

Interpreting Health and KPI Values

So you are looking at the main NMIS dashboard and you see that a node has a Health score of 92.2% as the example below, there is also a red arrow beside that, which is the result of the longstanding NMIS feature for auto baselining, this red arrow is pointing down, meaning that the health now is lower than the last period. So WHY is this node less healthy now than it was before, clicking on the node will reveal the KPI scores and you can start looking at what is changing.

You see this KPI summary again, you can see the overall breakdown of the health metric represented in the KPI values and you can see that the MEM KPI has a red arrow pointing down, the auto baselining is showing us that the Memory score is lower than previously with a score of 2.04 out of a possible score of 5. If we look at the graph for the last 2 days, we can see that the average value for the MEM KPI is 2.28%, showing us that the memory utilisation has increased a little.

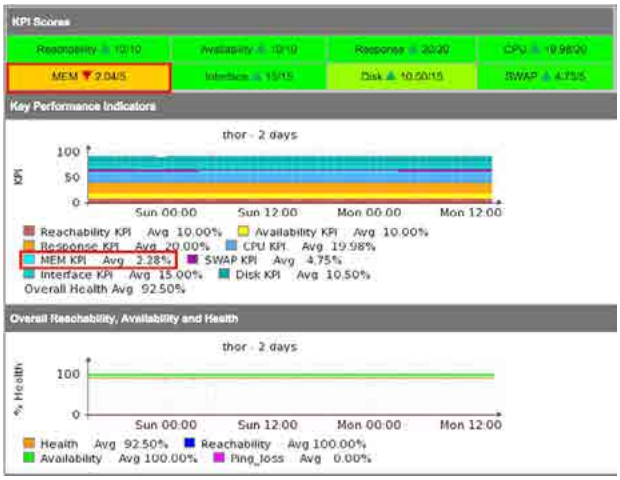

If you want to know WHY the health from the front page is 92.2% we can look at all the KPI values, like the Disk KPI of 10.50/15, CPU KPI is 19.98/20 and SWAP KPI is 4.75/5, we can take 100% and subtract the remainders so:

KPI ITEM	KPI SCORE/B>	REMAINDER CALCULATION	HEALTH REMAINDERS
Reachability	10/10	10 – 10	0
Availability	10/10	10 – 10	0
Response	20/20	20 – 20	0
CPU	19.98/20	20 + 19.98	0.02
Memory	2.04/5	5 + 2.04	2.96
Swap	4.75/5	5 + 4.75	0.25
Interface	15/15	15 – 15	0
Disk	10.5/15	15 – 10.5	4.5

Adding together the Health Reduction results and subtracting from 100 gives us: 100 – (0.02 + 2.96 + 0.25 + 4.5) = 92.27%

NOTE: The difference between the result and the displayed numbers are rounding precision.

Advanced Concepts

Analyzing this data will provide you with a lot of insight into your network behaviour. The metrics are designed to compare periods and work out if the current period (default 8 hours) is performing better or worse than the previous periods. We are comparing apples to apples in this scenario, but there is still deeper analysis that can be performed.

opTrend is a commercial module that can add a deeper layer of trending analytics to what NMIS offers, it will use 6 months worth of data to build a baseline, that is a true representation of a device, factoring into account the time of day and also the day of the week. This builds an accurate snapshot of what your system should be doing at a given time and assists in identifying outliers that aren't considered normal.

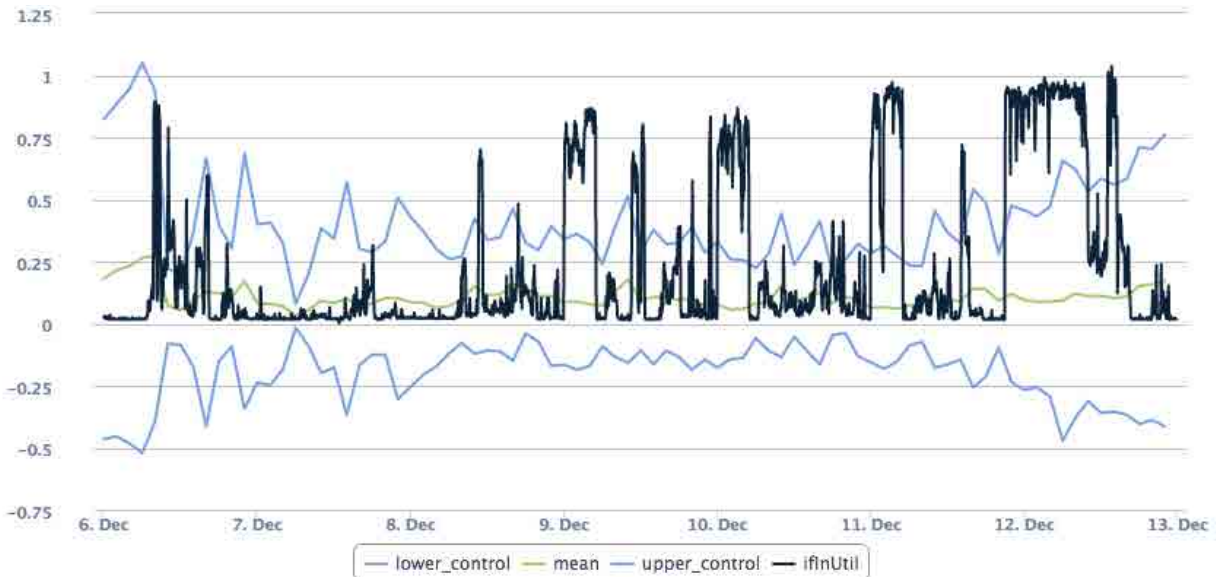

The graph above demonstrates the power of the application; the blue lines are the expected ranges, the green is the overall mean, and the black is the recorded value. The graph indicates that on December 6th, the system was in heavy use, but it was expected and not considered an outlier, this may be due to a scheduled system back-up for instance. However, from December 9-13 there was a significant increase that is not regarded as normal and would need to be investigated.

About the Author

Daniel Carter

Marketing Coordinator

Daniel Carter is a Marketing Coordinator with the Opmantek team located on the Gold Coast of Australia. Dan has an undergraduate degree in Psychology and a Masters degree in Information Technology. He joined the Opmantek team to increase the customer focus of the company. He is an adamant hockey and surfing fan.

Thank You

Need a Hand?

If you'd like some guidance around setup and configuration, or If there's a specific issue you're trying to solve in your network environment - click below to request a demo from an Opmantek engineer.

REQUEST A DEMO

OPMANTEK