

CUSTOMER STORY

Leibniz Supercomputing Centre

As part of the Gauss Centre for Supercomputing (GCS), the Leibniz Supercomputing Centre (LRZ) is one of the world's leading supercomputing centres with the high-performance computer SuperMUC-NG. The institute of the Bavarian Academy of Science and Humanities is focused on supporting IT services for science. At the same time, it concentrates on emerging technologies, Artificial Intelligence and Machine Learning as well as Quantum Computing in the field of Future Computing.

For more than 60 years, the LRZ has been a reliable IT partner for universities and scientific institutions in Bavaria, Germany and Europe. It provides a complete range of IT services and

technologies, as well as consulting and support – from email and web servers to internet access, virtual machines, cloud solutions, and the Munich Scientific Network (MWN).



Based in Garching near Munich, Germany, the Leibniz Supercomputing Centre (LRZ) is today one of the foremost European computing centres in scientific research by academic communities. LRZ supports ground-breaking research and education throughout a wide range of scientific disciplines by offering highly available, secure and energy-efficient services based on cutting edge IT technology.

The LRZ has been operating world-class supercomputers for decades. The current supercomputer, SuperMUC-NG, ranks among the most powerful computers in the world. With a peak performance of 26.9 Petaflops (almost 30 quadrillion operations per second), 719 Terabytes main memory, 50 Petabytes external data storage and a highspeed interconnect, the SuperMUC-NG provides first-class information technology for researchers

in the fields of e.g. physics, chemistry, life sciences, geography, climate research, and engineering. Funded by the federal and state governments with 125 million Euros, these computational resources are utilized throughout Germany, as well as for European collaborations. Last but certainly not least, SuperMUC-NG's innovative hot-water cooling system makes it one of the most energy efficient supercomputers worldwide.

The Challenge

”

I had a lot of fun with Icinga's own configuration language, which is very clever and practical.

Dr. Markus Michael Müller
System Administrator
High Performance Systems
Department
Leibniz Supercomputing Centre

Assuring Stable Operation of a Complex System

Responsible for Monitoring the high-performance computer SuperMUC-NG are Dr. Markus Michael Müller and Dr. Alexander Block.

Throughout the entire computing process, the LRZ focuses closely on supporting their users so they can take optimal advantage of all the resources the LRZ offers. When scientists are running their complex and resource-intensive simulations, the challenge is to assure stable operation of the intricate system, and to identify any issues before they turn into significant problems, without impacting the performance of the supercomputer.

Hierarchical Structure as Icinga's Crucial Benefit

When Markus Müller joined the LRZ in 2008, Nagios had been the monitoring tool of choice. However, this monolithic system showed limitations for their use case. Therefore, with the first SuperMUC system – the predecessor of today's SuperMUC-NG – LRZ migrated to Icinga 1, which already had a hierarchical system. In 2018, with their current leadership-class system, they upgraded to Icinga 2 and have been very satisfied with it ever since.



The Solution

High Availability Set up with Satellites

The system in place for the SuperMUC-NG consists of 2 Icinga 2 master servers in a high availability configuration and 36 Icinga 2 satellite servers. 7825 hosts and more than 76.000 Services are monitored.

SuperMUC-NG consists of 6480 compute servers that are connected via highspeed Omni-Path interconnect. The compute nodes are partitioned into 8 domains (islands). Within one island, the Omni-Path network topology is a "fat tree" for highly efficient communication. The Omni-Path connection between the islands is pruned (pruning factor 1:4). Each island accommodates around 800 compute servers and is monitored by 4 Icinga 2 satellites.

Soon their cluster will be extended: SuperMUC-NG Phase 2 will feature 240 accelerated compute nodes, which will be monitored by two additional Icinga 2 satellites.

Icinga 2 fulfils all their requirements in health monitoring and is very stable, never creating any problems. Markus Müller is also very satisfied with the Icinga 2 documentation that has helped him solve all problems without any support so far.

To create automated processes the team utilizes functions with preconditions. Markus Müller explains that " We use InfluxDB writer for performance data, Grafana to display trends in descriptive displays, which are then integrated back into Icinga Web.

This is, of course, a fantastic workflow."

”

For the LRZ, Icinga's built-in hierarchy is crucial, otherwise we would not be able to do monitoring.

Dr. Markus Michael Müller
System Administrator,
High Performance Systems
Department
Leibniz Supercomputing Centre

Ensuring proper Functioning of Hardware and Network

System administration of SuperMUC-NG is conducted through ethernet network, and they utilize around 200 ethernet switches that are closely monitored to guarantee uninterrupted access.

The hardware status of the compute nodes is monitored "out-of-band" through that ethernet network via the BMCs using IPMI and Redfish. In order to leave as much compute performance to the users' simulations as possible, checks which are running on the compute nodes are reduced to the very minimum and includes load, memory used, status for the batch system, and the Icinga 2 service itself.

Large simulations typically produce large amounts of data written to the filesystems. Of course, it must be prevented that the filesystems are filled up completely, because that would render them unusable. To this end, filling level and throughput are monitored and displayed, such that outdated data can be deleted as required. Beyond that, Icinga also monitors the proper functioning of the network and file system hardware. By utilizing IBM Spectrum Scale, the team receives alerts when a hard disk or disk shelf is malfunctioning.

Additionally, the cooling of their compute servers is done through hot water and is also under constant observation. According to Markus Müller "The role of Icinga in this context is highly significant, as it can promptly detect any pump failure and prevent the computer from overheating."

In addition to Icinga, they also use Splunk for logfile aggregation, which aids in analysis, but does not aid in issue prevention.

”

The role of Icinga in this context is highly significant, as it can promptly detect any pump failure and prevent the computer from overheating.

Dr. Markus Michael Müller
System Administrator,
High Performance Systems
Department
Leibniz Supercomputing Centre

Sharing Vital Information for All Departments

The outcomes of Icinga, represented through simple dashboards with color-coding, are shared with various departments.

For example, the LRZ CXS team, that supports SuperMUC-NG users, requires a quick overview of available filesystem space. Aggregated data on power consumption and cooling circuits are communicated to the facility management team. Furthermore, also the hardware vendor's support team depends on the information provided by Icinga.

While they already have had performance data in Icinga 1, it was not yet stored in a central database.

Markus Müller explains that "InfluxDB Writer opened up new possibilities for us. Icinga can write to the database and other departments' monitoring systems can then pull all necessary information from that backend DB. This way, we could remove needless redundancies in the monitoring."

An example for this is a data centre-wide monitoring system that focuses not on alerting the service quality of individual services, but rather on trending the vital statistics of the data centre itself, such as power and cooling.

”

What absolutely impresses me is the excellent documentation. Everything is documented fantastically, and I don't have to constantly read source code.

Dr. Markus Michael Müller
System Administrator,
High Performance Systems
Department
Leibniz Supercomputing Centre



Success



It's becoming increasingly difficult to convince me to use something else.

Dr. Markus Michael Müller
System Administrator,
High Performance Systems
Department
Leibniz Supercomputing Centre

Ready for the Future

And the next project with Icinga? Markus Müller and Alexander Block aim to implement Active Response for certain things to automate problem solving.

They also intend to make the collected data even better available to other departments to streamline the overall monitoring process and standardize previously separate solutions. Markus Müller sees Icinga as perfect solution and definitely wants to continue using it and even convince other departments to do so.

Outcomes

- ✓ Transforming Monitoring with High Availability Set up and Hierarchical Structure
- ✓ Assuring Stable Operation of a Complex System
- ✓ Sharing Vital Information for All Department
- ✓ Streamlining Overall Monitoring Process



Share your Story

Do you also have excellent experiences with Icinga and would like to [share](#) them?

We'd be happy to make your story come out big on our website!

Please get into contact with us at:

info@icinga.com

