Rittal – The System.

Faster – better – everywhere.

White paper – adaptive IT infrastructures

ENCLOSURES
POWER DISTRIBUTION
CLIMATE CONTROL
IT INFRASTRUCTURE
SOFTWARE & SERVICES

FRIEDHELM LOH GROUP

Contents

List of figures

Executive summary

Data centres and the networks between them form the backbone of today's information technology (IT). More and more services are being outsourced to the “cloud”. This has the advantage that the user’s terminal devices can be scaled down more and used longer by outsourcing computing work and storage space. On the other hand, this trend also places a higher burden on the networks and increasing demands on data centres. Despite the increasing energy efficiency of the hardware, data centre energy demands are growing, for one thing due to the increasing number of servers and storage systems. Every single Watt-hour of energy used by the IT must then be dissipated from the data centre in the form of heat, so that the energy demand rises again.

The AC4DC project (Adaptive Computing for Green Data Centers) [Ref. 1] aims to counteract the trend of rising energy demand in data centres by addressing the inefficient use of existing IT and infrastructure systems and managing them more efficiently. The most important technical objectives in the context of this paper were:

- Implementation of a load and power management system of services and servers for dynamically and proactively adapting the server capacity to the actual demand.

- Realisation of a comprehensive climate control for the dynamic optimisation of the entire cold chain, depending on the current IT performance.

- Optimisation of the operating costs in a data centre network by moving workloads to lower cost locations.

Due to the relatively wide spread of these technical objectives, in the AC4DC project, an almost complete data centre system was studied in terms of energy technology and optimised. Accordingly, the main objective was to combine the individual measures in an overall optimisation programme. While the load and power management ensures that the IT load behaves dynamically, comprehensive regulation attempts to adapt the climate control to the now dynamic heat load of the IT.

Introduction

The aim of the project AC4DC (Adaptive Computing for Green Data Centers) [Ref. 1] in the context of the IT2Green research initiative [Ref. 2] is to sustainably reduce energy consumption in data centres. To this end, intelligent forms of managing the computational load and infrastructure within a data centre (as well as across data centres) have been investigated in the context of the project.

Figure 1 illustrates the interplay of the individual sectors.

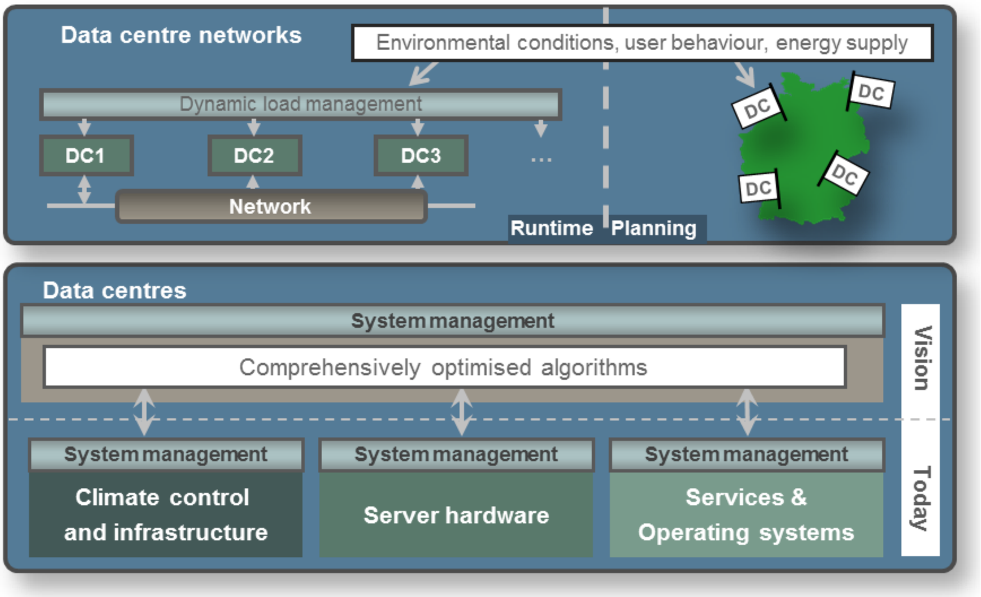

Figure 1: Comprehensive load management

In terms of the infrastructure of data centres, the focus is on the need to optimally adjust climate control across all levels. This not only includes the appropriate components and necessary architecture, but also higher-level control algorithms that adaptively adjust to demand.

In Figure 1, this is symbolised by the dark green “climate control and infrastructure” box, where the system management of the infrastructure has the task of a comprehensive regulation.

In a second step, communication with a higher-level load management system is added so that server management, services and IT infrastructure are comprehensively optimised.

Reference architecture

The modular RiMatrix S data centre was used as reference architecture [Ref. 3, Ref. 4, Ref. 5]. It has been designed with a regard to optimal climate control adjustment for small and medium-sized data centre solutions.

Server modules - as shown in Figure 2 - are available with or without an integrated UPS. These are already complete units that only have to be supplied with power, cooling (cooled water) and a network connection in order to operate the IT components (servers, data storage, and switches).

Figure 2: The “Single 9” data centre module of the RiMatrix S modular system

The “Single 9” data centre module in Figure 2 is found for eight server enclosures and one network enclosure, which are connected to the sub-distribution. The cold zone is formed in front of the server level. The hot zone – at the rear of the server enclosures – is separated from the cold zone by a partition.

Only six server enclosures are integrated if a UPS is installed. The UPS can be found together with the battery in a technically climate controlled and separated technology section. The separation of data centre and technology section allows the server air inlet temperature to be decoupled from the lower temperature in the technology section, so that two climate zones are created.

The raised floor of the server modules not only serves to improve the air routing but also to accommodate the climate control units for cooling distribution, as shown in Figure 3. The climate control units in the raised floor were moved to accommodate as many server racks as possible of within a server module.

The perforated raised floor panels possess an integrated high-performance fan, which blows cold air directly in front of the server racks. The heat exchanger, which cools the air using cold water, is located in the raised floor under the respective enclosure. The space above the enclosure suite is partitioned off, so that there is a cold zone in front of the servers and a hot zone behind them.

Figure 3: Air routing of the server modules

In order to perform cooling, the heat exchangers of the server modules have to be supplied with cooling water. The necessary refrigeration module consists of a large, V-shaped free cooler, redundant chillers (cold water units) and a pumping station at the head.

The RiMatrix S refrigeration module – see Figure 4 - is optimally designed for the RiMatrix S server modules and is available in two performance classes (70 kW and 100 kW).

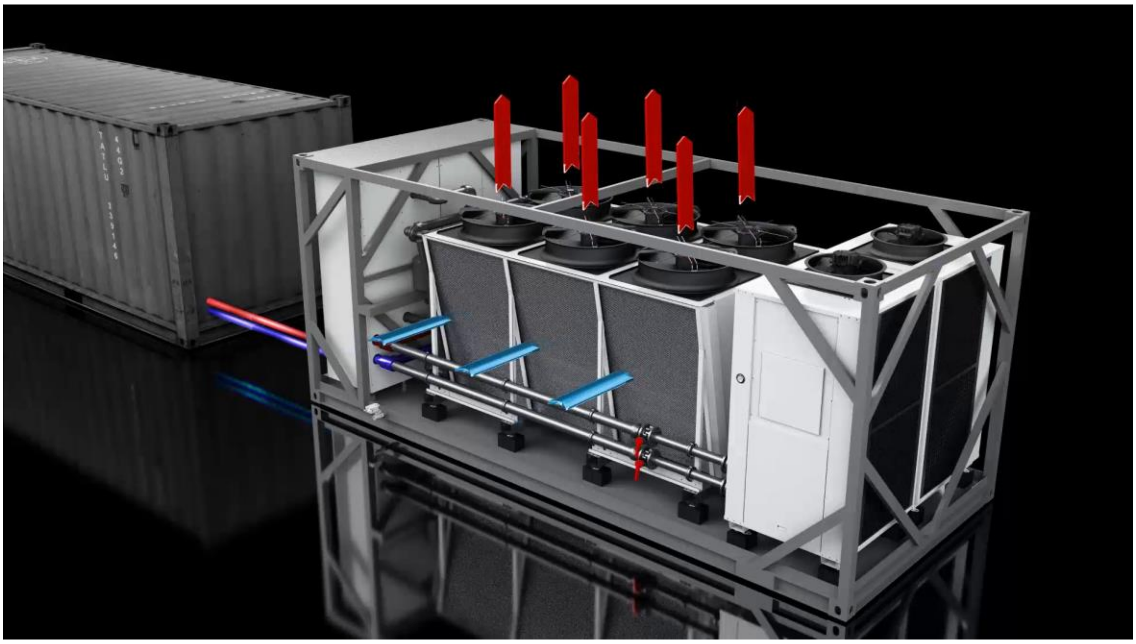

Figure 4: The cooling module

Objective of a comprehensive regulation system

First objective: optimal energy operating point

The minimum overall energy consumption for the data centre in a given time interval (t_1, t_2) must be found. The load setting of the active components is a constant to which the infrastructure has to be optimally adjusted.

$$
\begin{aligned}
E &= \int P(t)\,dt \\
&= \int (P_{IT}(t) + P_C(t) + P_{UPS}(t) + P_{FC}(t))\,dt
\end{aligned}
$$

The power P is derived from the power consumption of the individual sectors:

P_{IT}: = Power consumption of the active components

P_{UPS}: = internal consumption of the UPS systems (to buffer the active components)

The components of the climate control system are variable in terms of further consideration:

P_C: = power consumption of cooling distribution in the data centre

P_{FC}: = power consumption of cooling

It follows that the sum of the load capacity of the cooling ($P_C(t) + P_{FC}(t)$) must be assumed as being at a minimum during the period under observation.

Here, the variables are the:

- Server inlet temperature
- Inlet temperature
- Flow rate
- Volumetric air flow in the data centre
- Volumetric air flow during cooling
- Temperature difference – Delta T between the server inlet air and exhaust air

Under the following boundary conditions:

- Compliance with the ASHRAE requirements for temperature limits and humidity
- Maintenance of the minimum differential pressure - Delta P, between server inlet air and exhaust air

Second objective: adaptive adjustment to load changes

The regulation of the data centre climate control system has to be automatically adapted to changes in load. The biggest challenge is with sudden load changes such as those that occur when blade server systems are switched on and off. Figure 5 shows one example in the reference data centre where the load was sharply increased from 30 kW to 50 kW and then removed again.

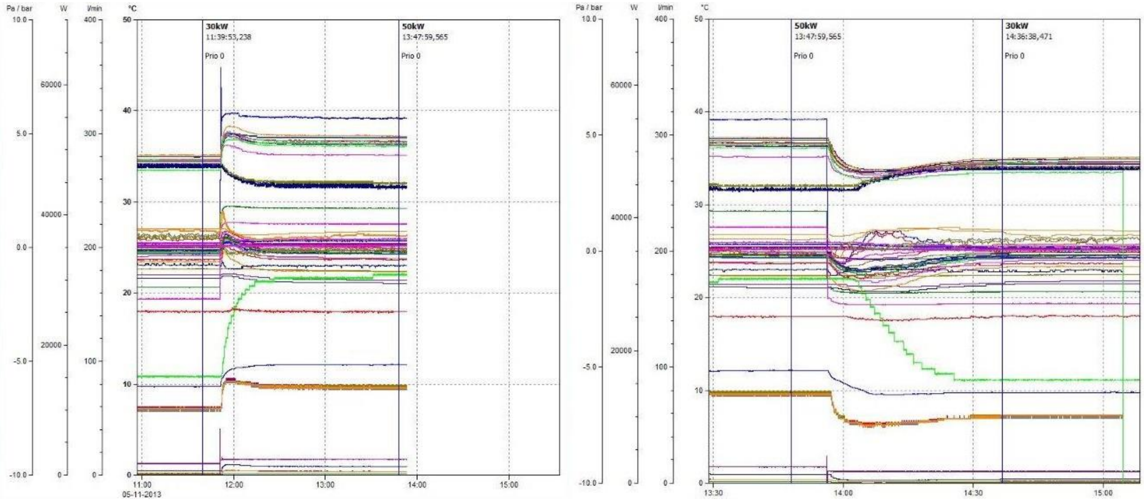

Figure 5: Adaptive load behaviour (30kW to 50kW, 50kW to 30kW)

Abruptly switching on a load leads to an increase in heat generation. For a certain period, the temperature rises in front of the server level, before it is again regulated to the reference variable. It is important that the temperature rise remains within the ASHRAE zone [Ref. 6].

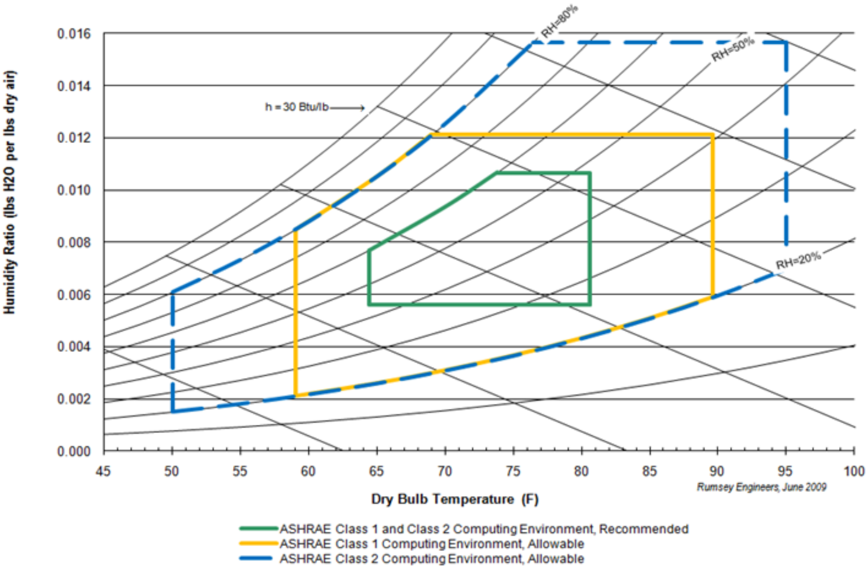

Figure 6: ASHRAE zone for temperature and humidity in the h,x diagram (© Ref. 7)

As the fans of the climate control devices respond quicker to load changes than the cooling system can provide additional cold water, a parallel shift between the heat rise in the data centre and the flow rate is observed.

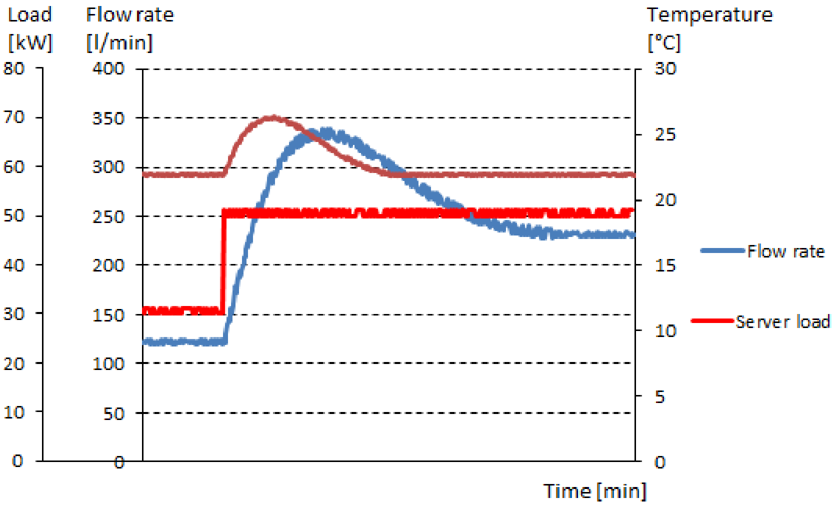

Boundray conditions:
Load: 30kW up to 50kW
Water-in temperature: 18°C
Temperature server-In: 22°C

Figure 7: Adaptive load behaviour (30 kW to 50 kW)

The course of the temperature and flow curves in Figure 5 and Figure 7 shows the behaviour of a damped oscillation, such as occurs in the control loop when using a PID controller, where the increase in server load triggers the controls (“jump request”).

Adaptive regulation

Delta-T regulation (ΔT)

The difference (ΔT) between the air inlet temperature and the exhaust air temperature of the IT equipment is critical for climate control. The higher the return air temperature is, the more energy efficient is the cooling. For the thermal energy Q, the following applies:

$Q := c\ m\ \Delta T$

where:

c := specific thermal coefficient

m_A / m_W : = mass of the transport medium (water / air)

ΔT_A / ΔT_W : = air-side or water-side Delta-T

Figure 8: Definition of the variables of a Delta-T regulation

That means the higher the air-side Delta-T (ΔT_A, shown in Figure 9), the less air must be recirculated to carry the waste heat away from the server. Likewise, the higher the water-side delta-T (ΔT_W) between water inlet and water return, the less water needs to be pumped between cooling generation and cooling distribution in order to transport heat from the data centre.

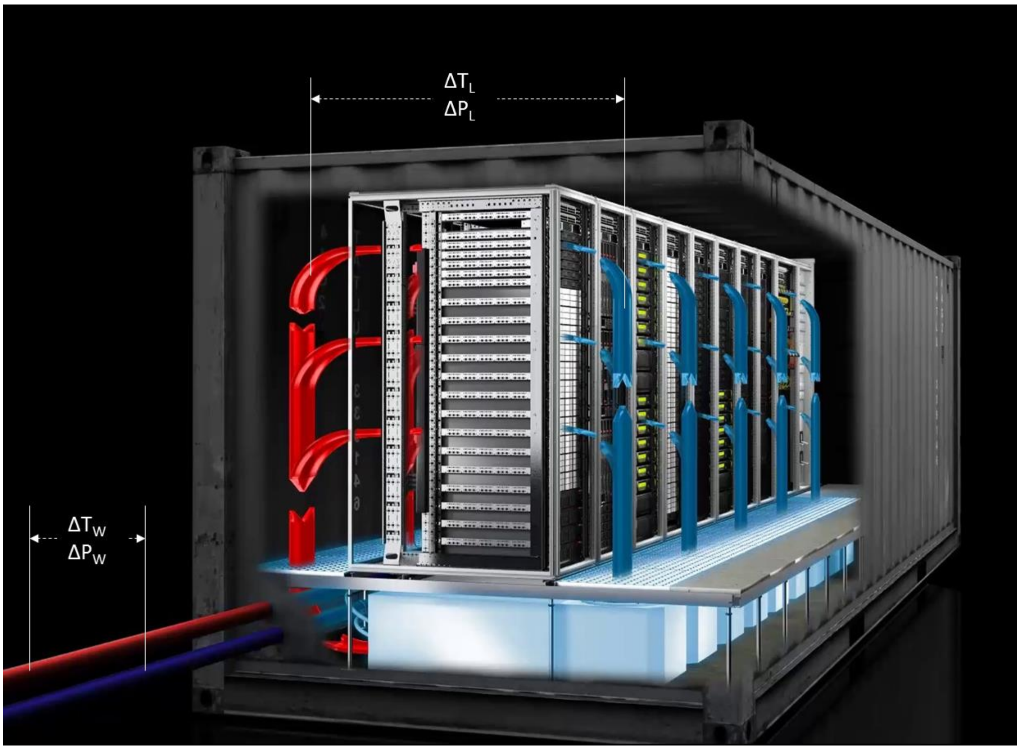

Figure 9: Water-side and air-side Delta-T and Delta-P

The ASHRAE conditions [Ref. 6] apply for the server inlet temperature, as well as for the humidity, so as to guarantee an optimum operating point for the IT components. In addition, Delta-T regulation must adhere in principle to the server-air volumetric flows.

With regard to cooling, the case of a constant, water-side Delta-T must also be considered. Attempts are made to keep the water-side Delta-T as large as possible, in order not to exceed a predetermined limit temperature in the return line. The cooling generation machines run very efficiently in this mode (facility mode).

Thus, the basic aspects of Delta-T control must also be taken into account in a comprehensive regulation system.

Delta-P regulation (Δp)

Here it is ensured that the servers are only provided with as much air as they can currently take in through their internal fans. The air with the lowest overpressure of approx. 1-2 Pa is provided on the cold side of the aisle. It should also be ensured that no local hot spots form that may arise from an insufficient volumetric air flow at too low a differential pressure.

A Delta-P regulation system (see Figure 9) ensures that the differential pressure between the inlet air and exhaust air of the IT equipment is constant and minimal.

Sliding inlet temperatures

The inlet temperature is automatically increased by the cooling system until a fixed, specified server air intake temperature can no longer be maintained. In addition, the external cooling system "may" for example independently change the flow rate, switch between the compression chiller and indirect free cooling or change the speed of the fans, compressors etc. However, the maximum server inlet temperature must be maintained according to the ASHRAE specifications [Ref. 6] under all conditions.

This way, an additional energy efficiency is achieved because the respective maximum inlet temperature increases the proportion of free cooling. The inlet temperature is dynamically tracked with the IT load. This means a significant increase in efficiency.

Comprehensive regulation

The aim of comprehensive regulation is to optimise all the components involved in the cooling process, taking into account the electrical power consumption. The control variable is the electrical power consumption of the entire system - this needs to be minimised.

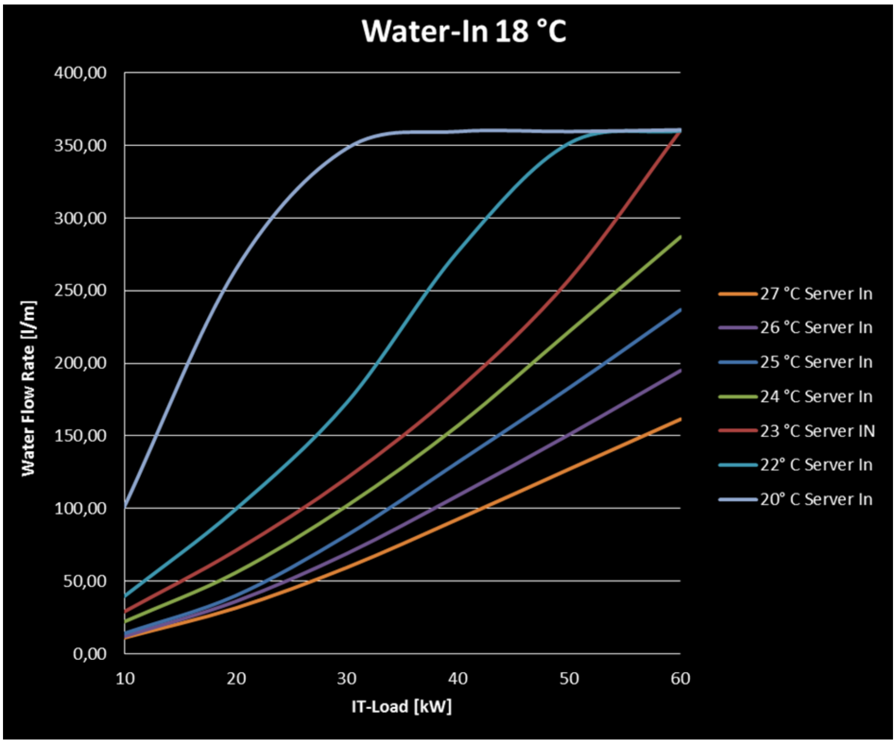

Figure 10: Reference architecture efficiency measures

Figure 10 shows an array family of curves in which the flow rate of the water is shown as a function of the IT load. At a given water inlet temperature, the various curves of the server inlet temperatures show how these result from the flow rate settings.

The previous control variable (server inlet air temperature) still applies, though only insofar as it represents an upper limit temperature.

The higher-level control system may change the parameters previously shown, as well as all the control variables of the climate control device (parameters of cooling distribution in the data centre). The influence of a “regulatory action” on the energy balance of the whole system is provided directly by measuring all the electrical different power consumptions in parallel. The individual parameters must be adjusted so that the minimum total energy is absorbed and the temperature limit is not exceeded.

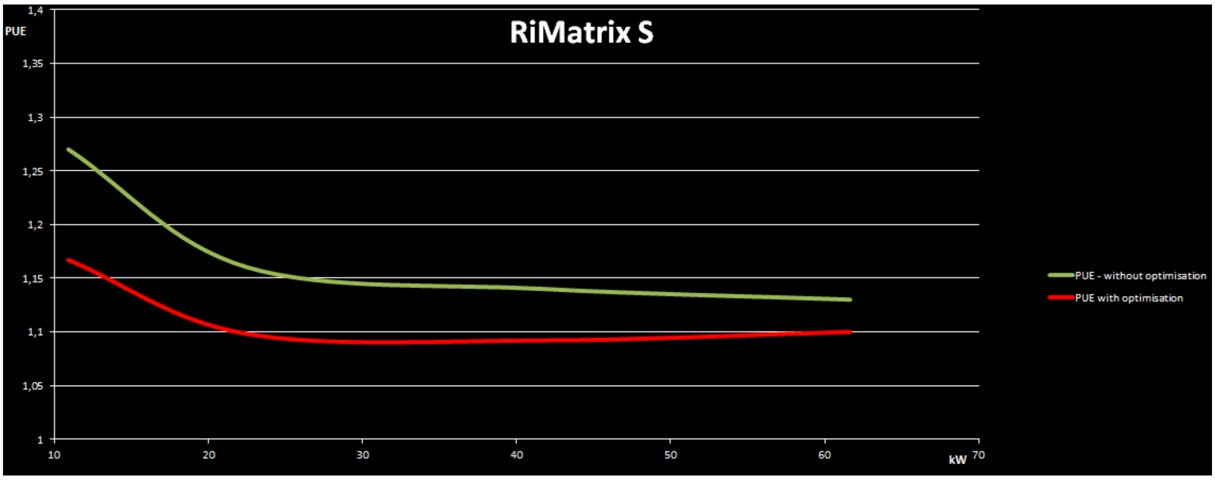

Figure 11: Results of comprehensive regulation

Figure 11 summarises the results of comprehensive regulation. A significant improvement in the efficiency (PUE) of the entire system can be observed.

It is ensured that the system consumes a minimal amount of electricity and is operated safely. If the inlet temperatures alone and thus the server inlet air temperatures were increased, the energy consumption of servers would rise very sharply from a temperature of approximately 27°C (depending on the servers used). Starting from a so-called “limit temperature”, the powerful fans of the server run at 100% of the velocity, which means a massive increase in the burden on the IT – the PUE is better, but not the overall energy balance. However, the data centre operator pays for the kWh and not the PUE. Therefore, comprehensive regulation that optimises the total power consumption makes sense.

Comprehensive load and power management

The RiZone Data Center Infrastructure Management Software - DCIM [Ref. 8] has an SNMP-based, bi-directional interface to the higher-level system management. The aim is to provide the server management (the higher-level Load and Power Management - LPM) with all the information it needs in order to achieve the optimal distribution of the load.

The communication interface described below is based on a multi-level hierarchy, as outlined in Figure 12:

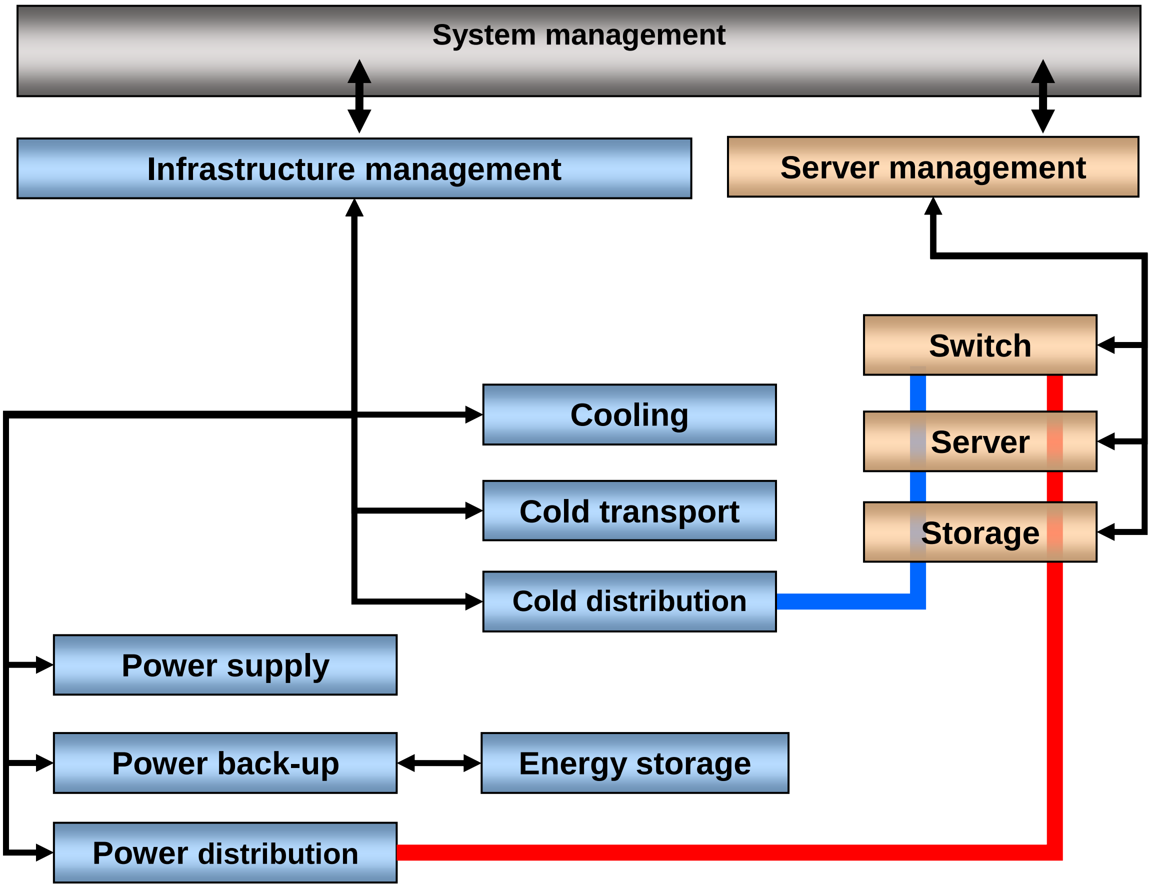

Figure 12: Communication with the system management

The most important communication parameters are:

- **The costs of cooling**

  Cooling costs are “fictitious”, uniquely calculated costs, dependent on the cooling variant selected (room, bayed, or rack climate control) in order to define preferred enclosures in a data centre where the load is preferably placed - possibly with different climate zones.

- **PUE**

  The efficiency of a data centre (the PUE, as well as the PUE values – hourly average value, PUE – daily average value)

- **IT energy consumption**

  Energy consumption of the IT (energy consumption of servers, storage, switches, etc.)

- **Energy consumption of the cooling**

  Energy consumption of the cooling (and the energy consumption of the free cooler, or the energy consumption of the chiller)

- **Temperatures:**

  External temperature (outdoors and where the cooling is located) server inlet temperature, exhaust temperature of the servers)

Thus, the LPM contains [Ref. 9] all the data essential for controlling, as well as for monitoring the data centre.

The objective of the dynamic load and power management of the services and servers is to adjust the number of pieces of active server hardware to the existing resource requirements of the services and their dynamic rearrangement during runtime. Services are defined by applications that need to be continuously operated and accessible.

As a prerequisite for applying the procedure, the servers to be managed must be operated in a virtualised environment and the services must be encapsulated in virtual machines (VMs). Through a technique called Live Migration, which allows the uninterruptible movement of VMs between servers, services can be managed dynamically and entire servers can be consolidated.

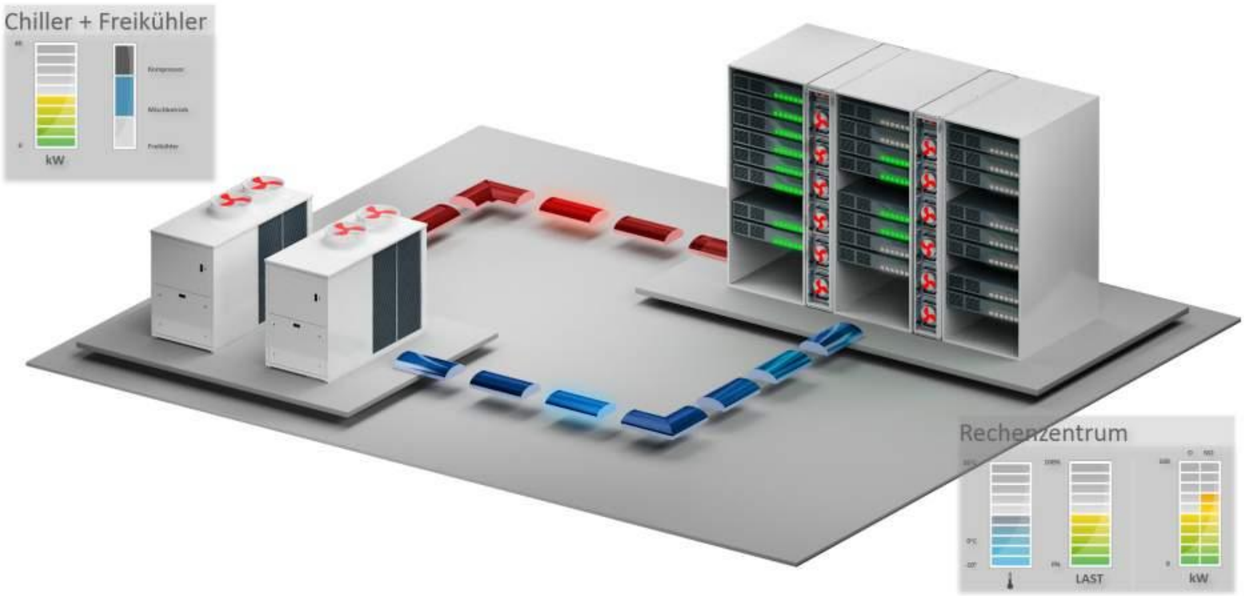

Figure 13: Interaction between server management with infrastructure

Comprehensive savings can thus be achieved with respect to the infrastructure. Thus, unnecessary servers can be switched off, as shown in Figure 13. However, the associated infrastructure is not required, either: Day/night or weekend operation with significantly different load profiles can thus be comprehensively achieved.

Literature

Ref. 1: AC4DC: Adaptive Computing for Green Data Centers; www.ac4dc.de

Ref. 2: IT2Green – funding initiative; www.it2green.de

Ref. 3: Rittal white paper: RiMatrix S - A strategy for designing standardised data centres

Ref. 4: Rittal white paper: The cooling technology of the RiMatrix S data centre

Ref. 5: Rittal white paper: The efficiency package of the RiMatrix S data centre

Ref. 6: https://www.ashrae.org

Ref. 7: https://nsidc.org/about/green-data-center/project.html

Ref. 8: Rittal white paper: RiZone - The Rittal management software for IT infrastructures

Ref. 9: Offis e.V. R&D Department Energy, www.offis.de

List of abbreviations

BSI	Bundesamt für Sicherheit in der Informationstechnik (German Federal Office for Information Security)
CMC	Computer Multi Control (sensor network system of a data centre)
CRAC	Computer Room Air Conditioner
DCIM	Data Centre Infrastructure Management
HPC	High Performance Computing
IT -	Information Technology
ITIL	IT Infrastructure Library (ITIL) is a specification for an IT service management system
LPM	Load and Power Management (server/application management)
LVDS	Low Voltage Distribution System
PDR	Power Distribution Rack
PDU	Power Distribution Unit
SNMP	Simple Network Management Protocol (communication protocol)
UPS	Uninterruptible Power Supply

Rittal – The System.

Faster – better – everywhere.

- Enclosures
- Power Distribution
- Climate Control
- IT Infrastructure
- Software & Services

RITTAL GmbH & Co. KG
Auf dem Stützelberg · D-35726 Herborn
Phone + 49(0)2772 505-0 · Fax + 49(0)2772 505-2319
E-Mail: info@rittal.de · www.rittal.de · www.rimatrix5.de

ENCLOSURES POWER DISTRIBUTION CLIMATE CONTROL IT INFRASTRUCTURE SOFTWARE & SERVICES

RITTAL

FRIEDHELM **LOH** GROUP