



SÜC – Highly available SAN for virtualised servers

No more downtime!

Service providers must guarantee high quality standards for your services. To guarantee the good level of their services,

SÜC Coburg relies on virtualisation and high failure safety, which is realised by a one-stop solution offered by N-TEC.

The Städtischen Überlandwerke Coburg, SÜC for short, must be available at all times for roughly 90,000 residents in the administrative district, no matter whether for electricity, gas, water, long-distance heating and bus connections. They constantly supply their users with the daily services that no one would, or could, do without. For that reason, the company knows how important it is to maintain this high degree of service quality. SÜC expects

nothing else from its EDP environment as well. Only a solid business and IT base will permit the fulfilment of the own customer standards.

SÜC has been an acting concern since 1999 and unites different offerings under one roof. These include the service departments SÜC Energie and H2O GmbH, SÜC Bus and Aquaria GmbH and the Städtische Werke Überlandwerke Coburg GmbH. The smooth operation of the

internal processes requires a solid EDP that comprehensively satisfies the various requirements. “We have a very mixed, heterogeneous IT environment”, explains Thomas Gruhn, EDP team leader of SÜC. “Add the fact that we also use virtualised servers. The memory must be reliable, scalable and powerful. We can only offer a good service if we have good working prerequisites”.

Serving server virtualisation and IT customers

Roughly 400 employees work within the SÜC group. A further 100 users were added because of an additional service sector. These are company customers who place their server services in the hands of SÜC. A dedicated line enables them to access their dedicated server, which could, for example, function as a domain server. The hardware and administrative expenditure is in the hands of the Coburg public utility company.

“These are the customers who rely one hundred percent on our services”, says Gruhn. “We cannot afford downtimes or bad maintenance here. That is how the requirement emerged about making the IT environment as reliable as possible, meaning highly available and flexible”.

Moreover, the group works with different servers, among others domain controllers, DNS servers, ERP systems, database and application servers. In the course of optimising the utilisation of resources, SÜC uses “VMware” “ESX” servers which generate a virtual server environment. However, this environment demands a high-capacity memory as regards not only the I/O load but also to the computer usage. The seven members of the EDP team headed by Thomas Gruhn looked for a corresponding Storage Area Network solution (SAN). This solution should not only be highly reliable but also expandable and guarantee optimal system utilisation. To this date, the ESX server had only used internal storage resources. It was not possible to either virtualise or allocate the capacities according to performance. Scalability or failure safety was not given either. The old hardware was to be replaced by a flexible environment in order to quickly provide new server services and be protected from failure risks.

Cluster SAN with ICEBOX

The choice fell on the N-TEC storage manufacturer from Ismaning, Germany. This company offered a clustered SAN with virtualisation function and the option of mirroring a location on a second location.

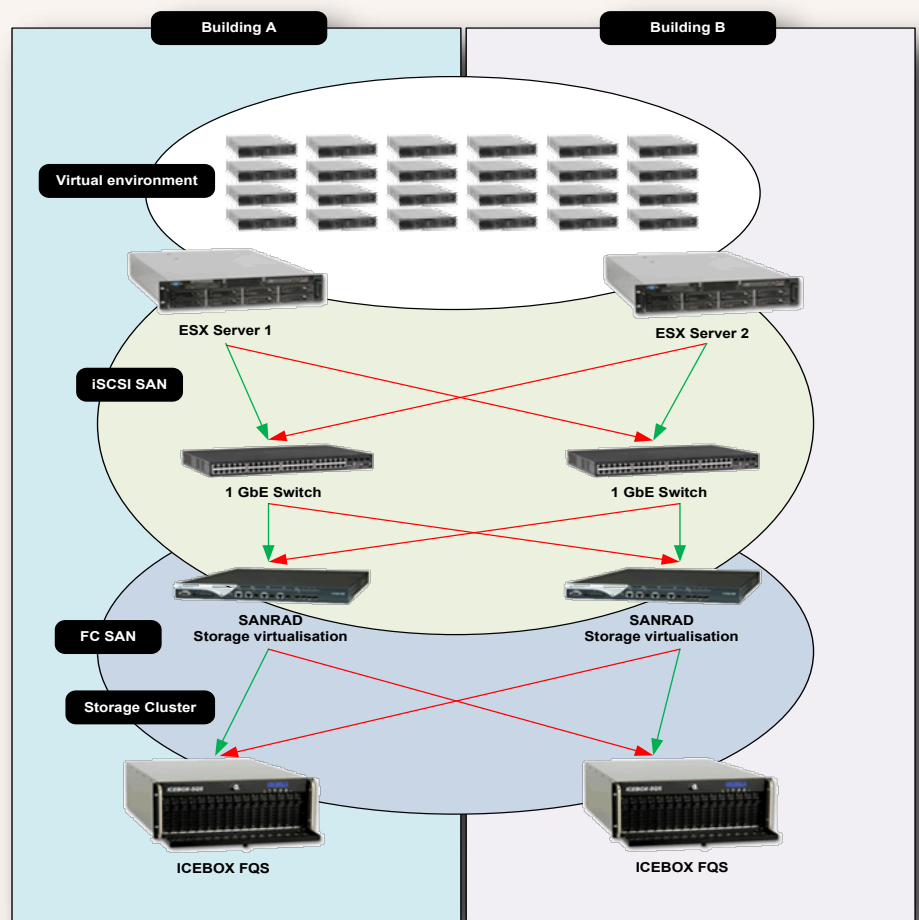
Two “ICEBOX FQS” RAID systems work together with two virtualisation switches: “SANRAD 3400”. In order to offer full redundancy and failure safety, these components were linked to form an iSCSI FC world in a dedicated architecture with two 1 Gbit/s switches and the ESX servers.

Each of the ESX servers are connected to both 1 Gbit/s Ethernet switches in the iSCSI environment. These switches in turn each connect individually to both SANRAD modules. A 4 Gbit/s fibre channel network helps the virtualisation entities find the connection to the ICEBOX FQS

RAID systems, which are also dual here: each switch to both storage devices. This crossed connection configuration alone already enables multipathing from the ESX servers. The failure safety also increases substantially. If one line is damaged, the entire data transfer can be “detoured” over the second route.

“This was the only configuration that was possible for these high requirements regarding failure safety,” explains Andreas Hauf, responsible for the project and the Head of Technology and Support at N-TEC.

“Redundant components, paths and data guarantee optimal protection. This helps SÜC to guarantee the best reliability even as a service provider.”



Goodbye failure: Dual paths offer high degree of safety.

To protect the data just as reliably, they are not just replicated asynchronously at SÜC as is done by many other solutions of other manufacturers. At SÜC, the data are synchronously mirrored for every write access to a second ICEBOX. The modified data blocks are simultaneously written to both ICEBOX systems. This is the only way to guarantee the redundant storage of application data which do not work on a release level (CIFS, SMB) but require access to the data medium in blocks in order to store the data. Data bases are the best example for such applications.

Pure replication solutions do not permit this. Moreover, the storage space of the ICEBOX systems can be efficiently used, or allocated, by the virtualisation switches. Gruhn can now allocate dedicated storage to individual virtual machines appropriate to use without being hindered by the limits of internal resources. Furthermore, memory from third-party providers can be integrated in the memory pool, thus offering SÜC greater play and flexibility.

“We are now able to offer the flexibility of our IT environment to our customers”, says Gruhn. “The virtualisation helps us to optimally use the existing memory. Also, the high redundancy and failure safety enables us to pass on this high degree of service reliability to our customers”.

Powerful, inexpensive and safe

Both RAID systems from the ICEBOX FQS model are the heart of the SAN solution. They offer room for 16 SAS hard disks with 147 Gbyte each and have two 4 Gbit/s FC interfaces which permit redundant data paths. The high-performance SAS disks and Intel “XScale 64 bit RISC” micro-processors ensure that the I/O load of the virtualised servers is transferred smoothly. The system permits mixed operation of SAS and

SATA disks, giving the group additional scaling options.

Additional storage space requirements can be realised by the SAS Expansion Port, which enables the connection of three additional JBODs. LUN addressing occurs independent of host. Dynamic RAID expansion, a so-called online expansion, and migration to a higher RAID level is possible. The ICEBOX also allows automatic rebuilds. The model supports Dual Loop as well as use in



Powerful & inexpensive Data is stored centrally on two N-TEC ICEBOX FQS systems

a Switched Fabric environment with SAN Masking.

To guarantee the already high degree of failure safety on the part of the system, N-TEC equips its models with redundant fans and power units. Also, the “S.M.A.R.T. Correction” of the manufacturer is useful. This technology includes a disk self-test, disk scrubbing and disk cloning. Scrubbing describes a data integrity check with recognition of bad sectors and the correction of parity errors. Cloning is the periodic examination of hard disks which swaps data to hot spare disks in case of errors.

The ICEBOX SAS Systems are administrated via the Web and are thus independent of operating system and platform. An additional software installation was not necessary for the EDP team. As a further configuration option, the system offers access via terminal emulation (RS232), SSH or the password-protected front control panel with display.

All system-critical components and parameters are monitored. In case of a fault, the IT manager is notified via SNMP Trap, email (SMTP), Win PopUp (net send) or by the network-independent Securion relay interface. The RAID controller is on a cold swap rack and can be exchanged quickly and easily as needed.

“The entire N-TEC solution not only offered us versatile function, sturdiness and safety,” says Gruhn. “The price-performance ratio was really

convincing. At first, the investment costs in the virtual environment appeared to be the same as that of new systems. But when considered in the long term, enormous savings potentials for the follow-up costs became clear.”

The SANRAD V switches represent the bridge to the IP world. They enable a complete storage virtualisation for multi-vendor FC storage systems. The switch permits memory allocation at a block level basis via algorithm which stores the data packages concurrently on both storage entities. Parity information is not stored, enabling the memory system to be directly connected to the server in case of an emergency. The data can still be used during this procedure, offering SÜC the possibility of adding further, other storage units.

One-stop solution from server to storage

The project that Thomas Gruhn and his team began already at the

end of 2007 first revolved around server virtualisation. They decided to use an N-TEC solution here as well. The “rapidServe” series models with QLogic iSCSI HBAs are used as ESX servers. These guarantee broad RAID support, good performance and compatibility.

“N-TEC gets bonus points especially for their good consultation and competence,” Gruhn affirms. “Their expertise in the server and storage environment enabled them to offer us

fore “anticipatory” for reading access, which reduces latency times.

Furthermore, the Ethernet switches can be exchanged without reducing the operating ability. Only the path – the iSCSI target – changes. The EDP team can add manufacturer-independent capacities for the storage, but the ICEBOX systems also allow internal and external expansions, enabling large storage volumes.



The central intelligence of highly available storage solutions is represented by both of the mirrored SANRAD 3400 switches

an individual, coordinated and highly functional solution. Virtual servers demand more than a conventional memory box. N-TEC was aware of the challenges of such an environment and specifically planned the implementation of a high-performance and failure safe storage network with us. The tests were successful so that we have the complete solution in productive use since August 2008”.

High scaling potential

Another criterion for choosing this solution was the investment protection which not only offers consolidation capabilities, but which also has an enormous scaling potential. The SANRAD switch is a switch optimised for virtualisation, enabling throughput rates of 180 Mbyte/s in a single 1 Gbit/s network. Equipped with three lines that makes up to 540 Mbyte/s. If converted to 10 Gbit/s (possible as of 2009), this would substantially increase the rate. These numbers are partially the result from the ratio of write-read accesses, which is 1:3. The cache works preemptively, there-

“The consultation and implementation of the project was done very professionally”, says Thomas Gruhn. “Service, price and scalability simply suited each other. We also have a high failure safety that we can use to satisfy the equally high service level agreements of our customers. The redundant data stock caused by real-time mirroring also protects us against data loss. Moreover, the great flexibility of this complete solutions saves us additional investment costs. This permits us as Stadtwerke Coburg to do justice to our reputation of being a reliable and quality true service provider”.

Imprint

Publisher: N-TEC GmbH,
Oskar-Messter-Strasse 14,
D-85737 Ismaning, Germany
Phone: +49 / 89 / 95 84 07-0
Fax: +49 / 89 / 95 84 07-11
Email: info@ntecgmbh.de
Web: www.ntecgmbh.de
Responsible according to the German Press Law Sven Meyerhofer
Editorial office: Ulrike Riess,
D-31700 Heuerßen, Germany
Layout: N-TEC GmbH