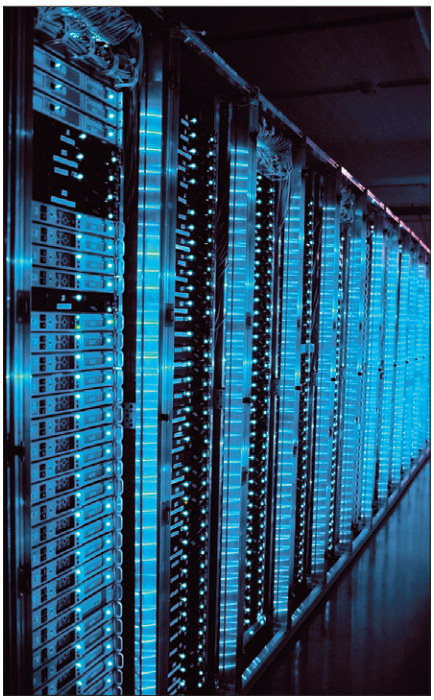

Figure 1: The 1&1 datacenter hosts around 40 percent of all domains registered in Germany – more than three million in total. Nearly all the 19" servers in a total of eleven computer rooms use Linux

Lots of Linux and impressive supply engineering

Linux Business

At the 1&1 datacenter in Karlsruhe, Germany, visitors can look forward to a mix of Linux machines and impressive supply engineering. The datacenter owner's primary mission – redundancy to the max. We took an exclusive tour of what is probably the largest Linux farm on earth.

BY MIRKO DÖLLE AND JAN KLEINERT

Visitors approaching the dazzling facade of 1&1's impressive six-story premises, and crossing a bridge over a landscaped lake to do so, will soon forget the cooked books of the dotcom bubble, and the rude awakening for many investors when that bubble burst. The architecture on this campus points to a simple, but nonetheless astounding fact: you can actually generate consistent turnover on and with the Internet.

After a short elevator trip down to the cellar of the 1&1 campus in Karlsruhe, Germany, and a complex clearance procedures to actually gain access, we finally reached the rooms that hold the machines behind the 1&1 success story: eleven computer rooms, each with no less than 62 19" racks containing various generations of Internet hosts – from shared hosting to root servers in single-height cases (see Figure 1).

Most of these machines run a Debian-based flavor of Linux, which was mainly authored by a team of 70 developers on-site. Six programmers are responsible for RAID controller drivers, or write patches designed to improve performance or kernel scalability. 1&1's Apache specialists have demonstrated their capability to support 15,000 parallel sites per single or dual-processor machine without impacting performance, whereas 1&1's competitors achieve only 50 percent of this figure at the best of times.

The team of 350 staff includes a hardware specialists who install about a thousand new machines in the racks every month. This kind of investment is absolutely essential: after all Netcraft's Internet statistics indicate that 1&1 is the fasting growing hosting enterprise with respect to server systems visible on the Internet – upward of 500 percent from March 2002 to March 2003.

The datacenter in Karlsruhe, which is run by 100-percent 1&1 subsidiary Schlund + Partner AG, hosts around 40 percent of all domains registered in Germany – more than three million – plus about 100,000 domains in the UK.

Figure 2: Jörg Hennig, the Chief Technical Officer of 1&1 subsidiary, Schlund+Partner, with the Cisco 1200 router to the DeCIX node in Frankfurt, Germany

Steffi and Martina Keep Things Organized

On what is presumably the world's largest Linux server farm, software distribution is managed by proprietary scripts. Individual servers are set up by distributing individual configuration files that reside in a Sybase database which itself is distributed across two Sun machines called *steffi* and *martina*. The database also stores customer data and logon credentials. 1&1's engineers use a highly-customized Netsaint variant for monitoring purposes.

Network Connection

Each server has a 100-Mbit/sec network connection; the individual computer rooms are in turn connected by 10-Gbit/sec FDDI lines. All of these converge in the router room, and from there data is routed across two 25-Gbit/sec FDDI lines (see Figure 2) to Paris and Frankfurt.

For 1&1, redundancy not only means twice the amount of machines, but also machines from different manufacturers. This helps 1&1 avoid a total breakdown of connectivity in case of manufacturer specific issues – such as the Cisco router software bug in mid-July of this year.

The Air-Conditioning Issue

A large collection of computers like this needs state-of-the-art air-conditioning and power-supply technologies. 1&1 invested no less than 13 million Euro in

power-supply, air-conditioning, and physical access control technology at its Karlsruhe datacenter. Under-floor ducts provide each 19" rack with a definable airflow, and each room is equipped with five or six air-conditioning units, each with a refrigerating capacity of 600 kW. The total airflow per room is somewhere in the region of 120 to 150,000 cubic meters per hour.

Redundant air-conditioning is important to 1&1: three air-conditioning units would be sufficient for emergency operations. All six units in each room have separate cooling water circuits, and each circuit is routed through every computer room. The cooling towers on the roof of the building work on a similar principle.

Mains Power and Diesels on the Roof

All the important utility services to the datacenter have been designed to be multiply redundant. The power-supply currently comprises of four independent lines with a capacity of 20 kV each, and line number five will be up and running shortly. These lines are normally serviced by the municipal power corporation based in Karlsruhe; they account for some three percent of the city's total power consumption. Emergency power is provided by a generator-based UPS system in the cellar, and a ready-to-run marine diesel engine with a capacity of 2 MVA is available for each line on the roof.

No End of Storage Batteries: The UPS System

The four power lines fed by the diesel generators lead to transformers that convert to 400 volt three-phase current and feed the UPS systems. At present four independent UPS systems rated at 1.1 MVA each are in service (see Figure 4); a UPS switching matrix allows their output to be fed into an arbitrary combination of power lines. The system also uses the generator matrix to compensate for loss of a complete power-supply line or transformer by selecting an alternative route to the generator.

Secondary voltage to the computer rooms is guaranteed by highly efficient synchronous alternators (95 to 97 percent) – it would be more or less impossible to cool a semi-conductor based UPS at these output levels. Dry lead acid storage batteries are kept in a separate room and are capable of providing backup power for 17 minutes at maximum output levels.

An additional switching matrix between the UPS systems and computer racks allows for uninterrupted re-distribution of UPS power between rows of racks. During normal operations, load is balanced between three UPS systems, the fourth one remaining on standby. The bottom level switching matrix allows operators to break the circuits to individual rows of racks in a worst-case scenario, thus ensuring the availability of emergency operations.

Where's the Control Room Operator?

All supply systems are monitored and controlled by computers. The complex air-conditioning system with hundreds of leak sensors and solenoids uses a control program more typically found in the chemical industry, where fluid systems of similar complexity are more common. The power-supply monitoring program is normally more at home in energy management, but traditional power management systems would be completely overtaxed if faced with 1&1's systems. Unfortunately, both programs are available for Windows only.

The monitoring software running in the control room is designed to handle numerous error conditions autonomously. In all other cases engineers are alerted by SMS, pager or phone. After all, it is their responsibility to cater for this huge penguin colony's every need. ■

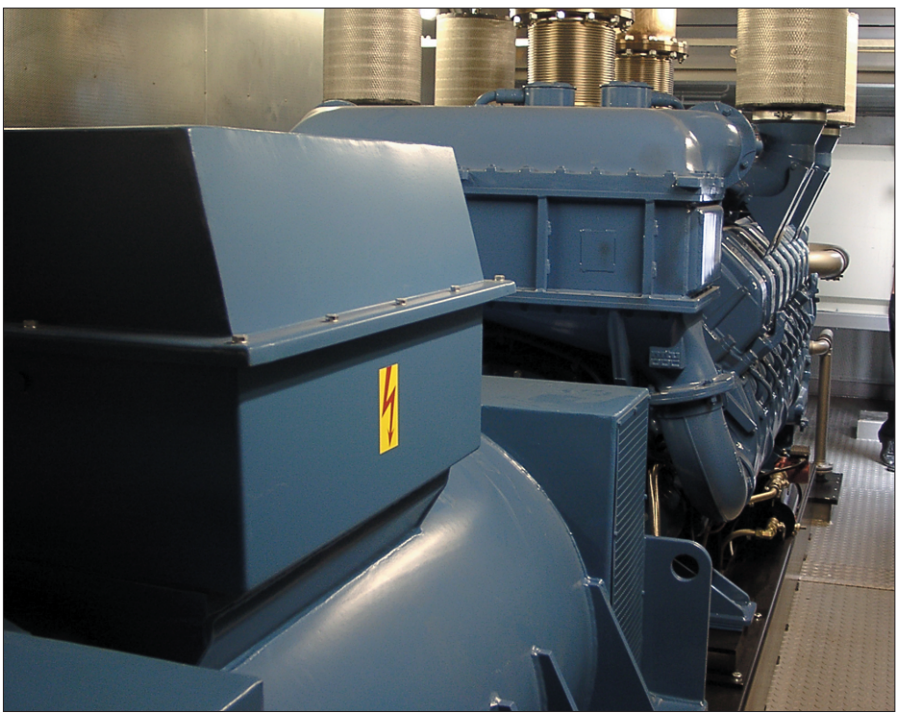

Figure 3: Four 16-cylinder marine diesels weighing 40 tons each have been installed on the roof of the building, and are used to supply the 20 kV internal power lines in case of mains power-failure

Figure 4: Each power-line is backed up by a 1.1 MVA UPS