Running on adrenaline.

WHEN AN ICT INCIDENT OCCURS, THERE'S NO TIME TO LOSE. EVERY SECOND COUNTS. TO ENSURE THE QUALITY MANAGEMENT TEAM CAN RESOLVE PROBLEMS AND IDENTIFY ROOT CAUSES FAST, T-SYSTEMS REGULARLY PUTS ITS EMERGENCY PROCEDURES TO THE TEST. EVERY YEAR, THE ICT PROVIDER PERFORMS OVER 500 "FIRE DRILLS", CAREFULLY OBSERVING HOW ITS PEOPLE, PROCESSES AND PLATFORMS INTERACT. WE LOOK AT HOW BOTH REAL AND SIMULATED INCIDENTS ARE STREAMLINING AND ENHANCING INCIDENT RESPONSE AND RESOLUTION.

*<**Copy**> Klaus Rathje*

Photo: Thorsten Futh

MANAGING QUALITY, ENSURING SUCCESS

T-Systems' holistic Zero Outage program ensures maximum quality and reliability in IT operations and transition and transformation projects. It includes the ongoing monitoring and fine-tuning of all systems and services plus continuous process development. A clearly defined 3P approach (processes, people and platform) enhances quality across all levels of the enterprise – with a single, consistent goal: achieving the best possible outcome for the customer.

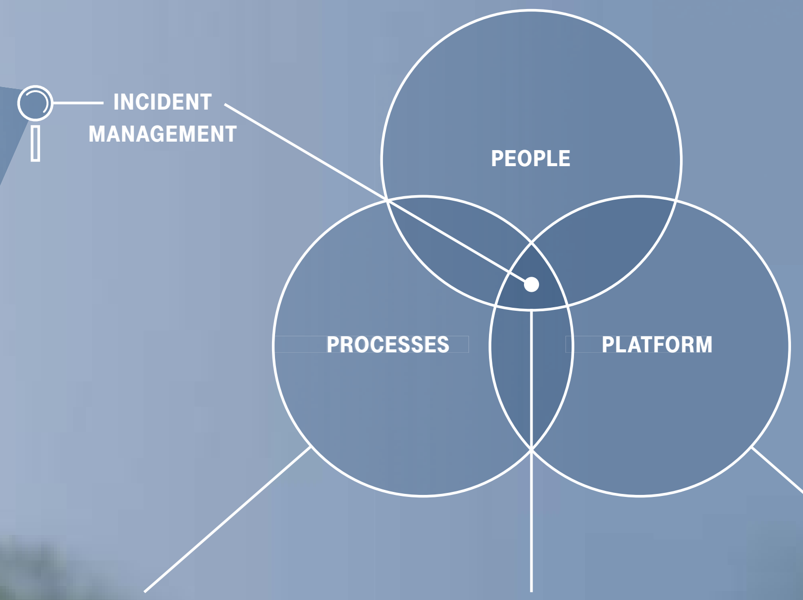

ADHERENCE TO CLEARLY DEFINED PROCESSES ENSURES FAST, TRANSPARENT RESULTS AND RAPID RESPONSES.

ADHERENCE TO PROTOCOL

Regular incident simulations improve response times and streamline emergency procedures. What's more, these training exercises ensure all stakeholders act in accordance with standardized processes. Weekly reports are turned into concrete proposals for improvements.

DUAL CONTROL

A strict two-person policy in critical cases eliminates faults and ensures realiable, effective quality control.

PEOPLE POWER – THE ZERO OUTAGE CULTURE IS BACKED UP BY INTENSIVE TRAINING.

QUALITY ACADEMY

The T-Systems Quality Academy certifies over 21,000 employees annually, helping them internalize the Zero Outage philosophy, culture and processes. This certification is renewed each year.

TOP-LEVEL MANAGEMENT ATTENTION

During weekly reviews, senior managers are briefed on processes, critical projects, disruptions and incidents. Involving top-level employees from across the organization can make all the difference, paving the way for faster decision making. If a major incident occurs, directors and senior vice presidents are brought on board. This approach ensures compliance with quality KPIs and effective, continuous improvement.

PROACTIVE CUSTOMER COMMUNICATIONS

A service interface consolidates customer feedback with the aim of enhancing processes and ensuring compliance with all business requirements.

A STRONG TECHNOLOGY BACKBONE IS THE KEY TO SUCCESS.

HIGH-AVAILABILITY IT

Reliable, redundant IT infrastructures hosted in twin-core data centers ensure maximum uptime for system landscapes. Comprehensive monitoring and regular testing of all components help quickly identify system overloads and faults.

SUPPLIER CERTIFICATION

Extending Zero Outage certification to suppliers ensures end-to-end IT availability across incident, problem and change management processes.

The Global Lead Incident Manager performs an initial assessment to decide which specialists Steffen Germersdorf needs to call.

AS STEFFEN GERMERSDORF KNOWS ONLY TOO WELL, being on call means having to drop everything at a moment's notice. That's why, even though today is a public holiday in Germany, he has not taken his eyes and ears off his phone. The 31-year-old works in Global MoD Service, T-Systems' international team for critical IT incidents. Global Lead Incident Managers like Germersdorf coordinate a worldwide pool of 1,000 experts, from a range of disciplines, who are available 24/7. Each week, 140 Managers on Duty (MoD) oversee incident resolution.

BERLIN, 10:57 AM:
FROM PIZZA TO PC

Germersdorf is something of an IT firefighter: "Sometimes, you'll work a shift and nothing happens at all. But as soon as the alarm sounds, it's all systems go." Only instead of battling blazes, Germersdorf and his team are combatting bits and bytes. "As a global ICT service provider, we typically expect to handle two or three critical events per shift," he explains. As the pool coordinator, his task is to identify the best people for the task at hand. On this particular day in May, Germersdorf is at home in Berlin, about to tuck into a slice of pizza when his phone rings. Before he has made it out of the kitchen and to his desk, the guy at the other end of the line has brought him up to speed: there's a problem in South-East Asia, six time zones ahead of Germany.

SINGAPORE, 4:57 PM:
60,000 USERS IMPACTED

It's a major headache: a WAN outage in Asia Pacific. Germersdorf has exactly 20 minutes to put together a team, and to set up a conference call between the customer service people, the IT delivery experts, and the software engineers. From now on, he is in charge of coordinating all troubleshooting activities. Full responsibility lies with him – and the clock is ticking. Each of the departments involved allocates a Manager on Duty to assist him.

The Global Incident Control Center that fielded the customer's call makes it crystal-clear that the situation is serious. "From the very first second, you can sense the adrenaline flowing," says Germersdorf. "WAN connection in Asia Pacific is down. Customer's production processes are disrupted; approximately 60,000 users affected." 60,000 staff with no network connectivity is a total disaster from the customer's perspective, and Germersdorf has to put himself in the customer's shoes – to gain insight into the incident, and to understand its impact on the business. Since the alert came in, half a dozen trouble tickets have been opened, but No. 8760047 gives the full picture – and confirms the worst: "Email communication, web traffic, financial processes, supply chain, operational processes, HR and production are all affected."

Photos: Thorsten Futh

INCIDENT MANAGEMENT PROCESS

All incident management activities follow a clearly defined process independently certified by TÜV Rheinland. Up to 500 training drills each year put the global teams through their paces.

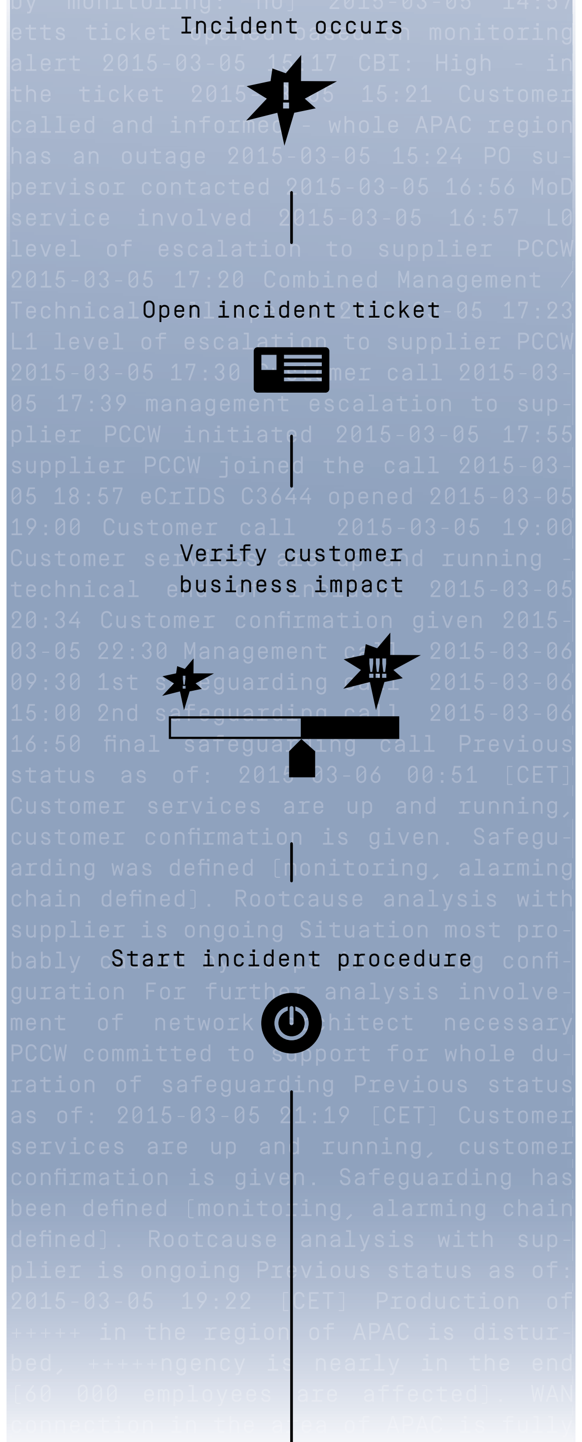

The Lead Incident Manager receives regular status alerts from the technical team throughout the process.

BERLIN, 11:22 AM:
INCIDENT TEAM READY IN RECORD TIME

During large-scale potentially critical incidents or events classed as critical by the customer, more than 20 experts attend these phone pow-wows, helping to streamline communications throughout the resolution process. A Global Lead Incident Manager coordinates the activities of participants from across the globe in a virtual war room. In addition to T-Systems employees, they include representatives of third-party vendors and partners – Cisco, SAP, Microsoft, VMware, F5 and Juniper, for example. And if the situation is especially complex, a T-Systems senior executive joins the discussion – because high-level incidents require the attention of high-level management. Germersdorf and his team have direct access to their technology partners' data and a designated contact, saving valuable minutes. Their goal is to comply with the mean time to repair defined in the customer's SLA.

Germersdorf's virtual conference is the key channel of communication and collaboration, keeping participants connected at all times. Every half hour, he delivers a status update, stopping only when the customer's IT is back up and running and the incident is resolved. Until then, everyone remains in the call. As Germersdorf points out, "passing the buck back and forth is the last thing you want in a situation like this."

All activities are based on a clearly defined process that has been independently verified by German testing and certification organization TÜV Rheinland. Every MoD knows the ropes, and has put their knowledge to the test many times. "Regular training keeps you on your toes," says Germersdorf. Against this background, T-Systems runs around 500 fire drills every year. These exercises simulate a real-world incident and involve the entire company, from call-center agents to top executives. Germersdorf explains, "Even if you're about to sit down to lunch or are celebrating Christmas with the family, if you're on call, you have to drop everything and respond." The first step is to establish the facts, and communicate them to the customer, the account manager and the engineers. You have to act swiftly, but in a structured way, in line with defined process. What's more, there are clearly defined KPIs in place to measure and record the success, or failure, of each step.

MAGDEBURG, 11:39 AM:
IS THERE A PROVEN SOLUTION?

T-Systems' control center for critical incidents is in Magdeburg. And whenever an alert comes in, the first thing anyone does is to take a look at past incidents. Have we seen anything like this before? Is there a tried-and-tested solution? If there is no precedent, activities follow standardized procedures. "We check which sites are affected and determine the worst-case scenario," explains Germersdorf. "This is the job of our technical guys."

SINGAPORE, 6:01 PM:
LOOKING FOR THE ROOT CAUSE

While the management call is in full swing, T-Systems' engineers run a series of checks on the customer's landscape. The news that the WAN is down does not narrow the search. On the contrary: servers, operating systems, applications, database operations, networks and data center power and cooling systems must all be looked at. Germersdorf reports: "While we're in the call, there are people trawling through multiple IT components." But that doesn't mean the pressure is off: his laptop is still overflowing with status alerts: "– cn-cti-shang-cx-09 ping statistics – 5 packets transmitted, 0 received, 100% packet loss, time 4022ms."

And while all this is going on, it is essential to keep the customer in the loop: "We would like to notify you that we are escalating this incident to our local partner. Our investigation is still ongoing. Unfortunately, no details are available on the expected system recovery time. We will be in touch as soon as we have further information. Best regards, Steffen." Providing regular updates and inviting the customer to play an active role in risk assessment and key decisions are top priorities throughout. "For example, if we need to reboot the system, it is the customer who decides when it should be done."

At this juncture, analysis has indicated that the matter is more complex than originally thought. But each additional step sheds new light on the situation. It would appear that, this time around, the cause is 'merely' a disrupted WAN connection. But Germersdorf is very aware of the consequences for the 60,000 employees affected. Until 2008, he was a member of the Magdeburg Global Incident Control Team. For the past seven years, Germersdorf, who holds a degree in IT, has worked in Global MoD Service. Today, he and his virtual team are working flat-out, trying to identify the root cause of this problem, which could be anywhere in China, Malaysia and Australia – or somewhere in-between. One by one, minute by minute, they are able to rule something out. Every time a potential cause has been excluded, someone on the team sends Germersdorf a message reading 'closed'. But then there is a real breakthrough: they manage to track the problem down to a number of network nodes. Three quarters of an hour later and the team is able to pinpoint two routers as the likely origin. One or both have clearly failed. "Network nodes are regarded as 'medium risk'", explains Germersdorf, in his first breather since he set to work several hours ago. "But when multiple medium-risk components fail simultaneously, you end up with a highly critical incident. Five or six isolated fires can fast become a blazing inferno."

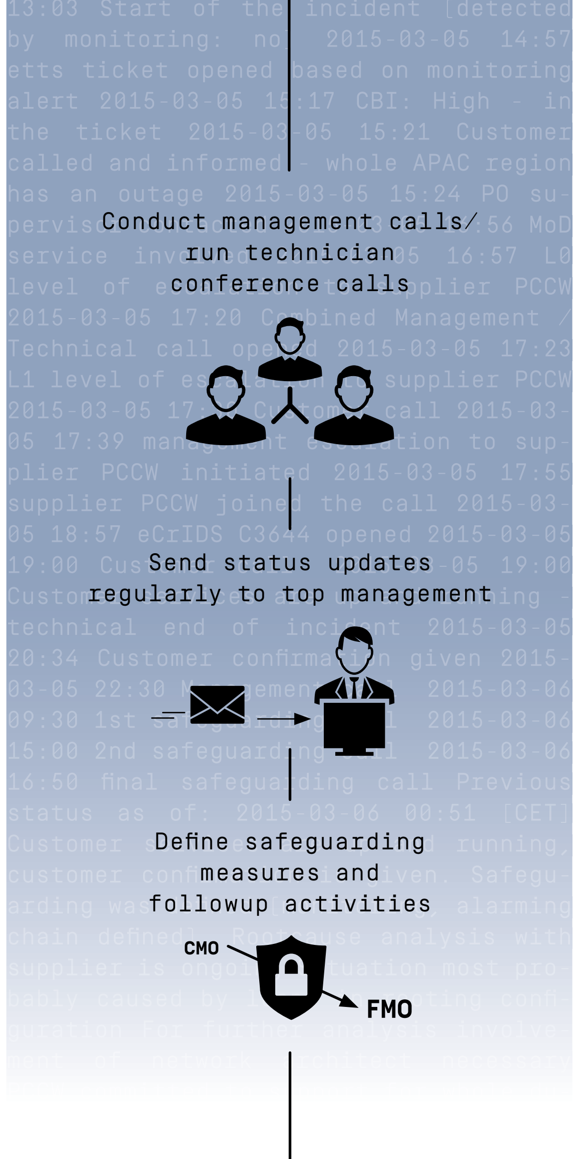
Conduct management calls/ run technician conference calls
Send status updates regularly to top management
Define safeguarding measures and followup activities
CMO
FMO

Photos: Thorsten Futh

CANBERRA, 9:50 PM:
A DIAGNOSIS IS SENT TO THE LOCAL SERVICE PARTNER

The incident management process, spanning eight time zones, is complete. The underlying problem was a router at a customer site in Canberra, Australia. The local service partner was quickly put in the picture, and was able to fix the issue within the hour. A member of Germersdorf's team relays her findings: "We discovered that an incorrect configuration command had been entered into the router and this was causing the connection to other autonomous systems to fail. As a result, the network node was down and access denied to 60,000 employees in the Asia Pacific region, seriously impacting communications." A tiny piece of grit on the optical fiber connection between routers can trigger an unexpected hardware fault like this.

Later, as part of the follow-up process, the router manufacturer will take an in-depth look at the configuration software, and provide the necessary patch. But to minimize the impact of such incidents in the future, the Lead Problem Manager, who Germersdorf has tasked with compiling a detailed report, has a suggestion: the service level defined in the customer's initial architecture design is at odds with the one that T-Systems deems necessary. To this end, Germersdorf's team will propose that the customer upgrades its medium risk network nodes to critical.

MAJOR INCIDENT = HIGH MANAGEMENT ATTENTION

"At T-Systems, major incidents require the attention of senior management. This is what makes us stand out from our competitors. All of our executives and Senior Vice Presidents receive around-the-clock status updates. This approach influences everyone in our organization, and is part of our DNA. What's more, it makes us focus much more closely on our customers. Each Senior Vice President is a manager-on-duty, and all board members are kept up to date. In highly critical situations, customers have direct, instant access to the T-Systems Board of Management – even on a weekend. This is Zero Distance in action – and in 28 years in the IT industry, I've not seen anything like this from any other provider."

Carsten Gram,
Senior Vice President, Big Deal Management

All of the customer's configuration items are monitored automatically, across the globe.

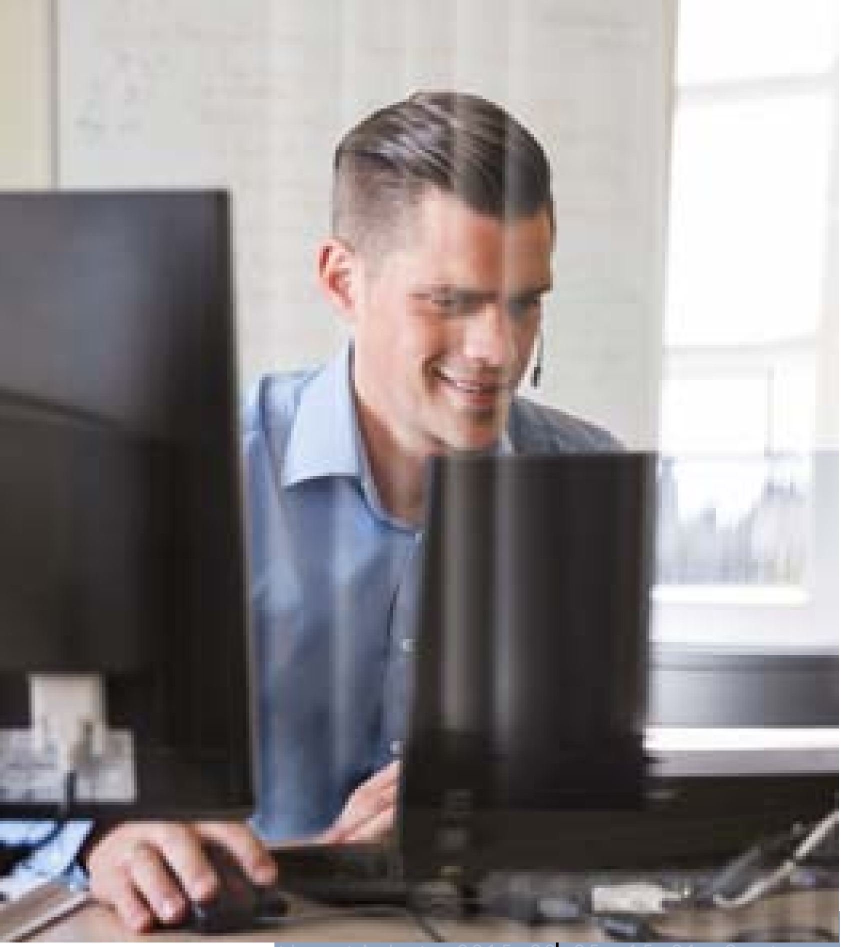

A sigh of relief – only when the customer's services are up and running can the team take a moment to catch their breath.

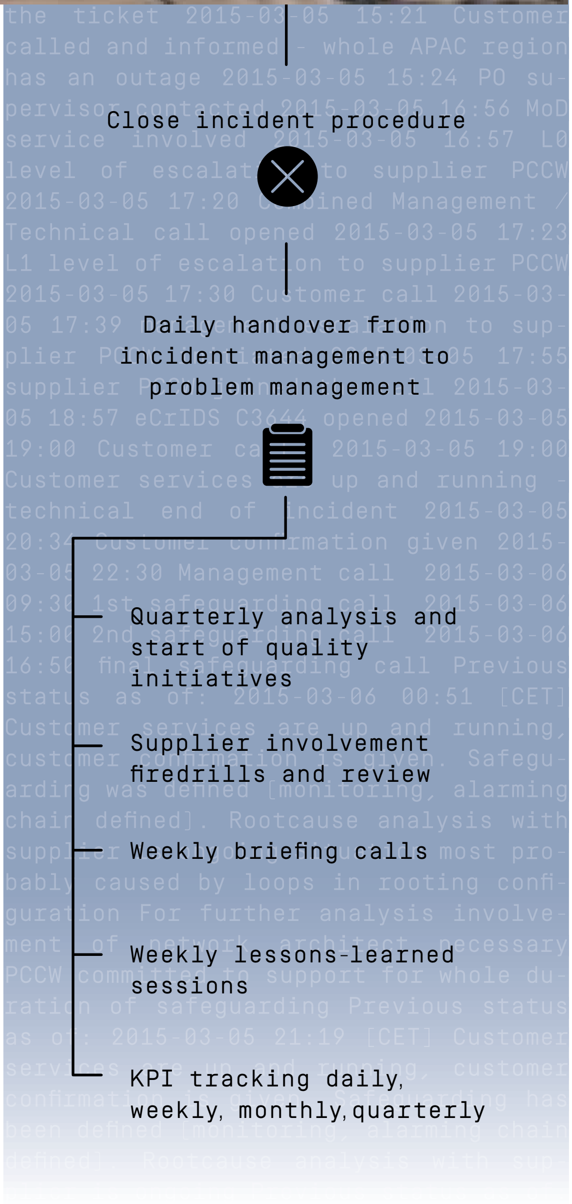

Close incident procedure

Daily handover from incident management to problem management

Quarterly analysis and start of quality initiatives

Supplier involvement firedrills and review

Weekly briefing calls

Weekly lessons-learned sessions

KPI tracking daily, weekly, monthly, quarterly

BERLIN, 4:51 PM:

THERE'S ALWAYS A PLAN B

Some events are so unexpected that you simply cannot prepare for them. At times like these, Germersdorf's advice is "just improvise". These unforeseen factors can range in size from a piece of grit like in Australia today, to 2000-meter-high volcanic mountain ranges such as Esja, Katla and Eyjafjallajökull in Iceland – that can be additional hurdles to incident resolution.

And this is exactly what happened last year when Icelandic volcano Bardarbunga brought European air traffic to a standstill. A T-Systems customer in Sweden was experiencing downtime, but the Global MoD Service team was on the case and knew how to solve the issue within a single day. A replacement server component was found in the Netherlands, and an engineer dispatched to collect it. But then came Bardarbunga, and the employee was left stranded at Amsterdam airport with the all-important part in his luggage. "It soon became clear that no aircraft would be taking off in Europe anytime soon," recalls Germersdorf. "So we quickly organized customs documents and had the part shipped by road and sea to Stockholm." In another instance, the situation was so urgent that a software expert was flown in by helicopter. But the extra effort and expense was worth it and the problem was quickly resolved.

"We try to take something away from every incident," says Stephan Kasulke, T-Systems Senior Vice President of Global Quality, "with a view to continuously reducing downtime for our customers" (see interview on page 31). The lessons learned from real-life incidents and fire drills not only help fine-tune crisis-management processes, they also help to predict potential IT incidents. Take configuration items, for example. "We monitor all CIs, including IT components such as servers, software and networks, very closely," explains Germersdorf. "This means we generally gain timely insight into any customer applications that are not running as they should. And we can see if the problem impacts one server that runs systems for 50 customers, or five servers deployed for just a single company."

BERLIN, 5:10 PM:

END OF INCIDENT

One more email pings into Germersdorf's inbox, and he breathes a sigh of relief: "End of Incident". The team has met the SLA target for mean time to repair, and all 60,000 users are back online. For Steffen Germersdorf the case is closed. He gets up from his desk to pick up where he left off in the kitchen. Luckily, the IT expert is prepared for all eventualities, and he has a back-up pizza in his freezer. After all his hard work, he has certainly earned it.

<Contact> *stephan.kasulke@t-systems.com*
<Link> *www.t-systems.com/zero-outage*

Photos: Thorsten Futh, Deutsche Telekom AG/T-Systems