

# Site Reliability Engineering:

\_AI-generated image. Prompts: helmet with attached cockpit instruments as sculpture, digital art.



## the autopilot in the cloud

### Expert article

\_Daniel Zeiter, Head of Technology

\_Laura Graf, Consultant

Ergon Informatik

Published in SMART insights 2023 magazine

**ergon**

smart  
people –  
smart  
software®

Companies are using Site Reliability Engineering to automate their IT operations. It's scalable (SRE) and dependable, allowing a fast and flexible response to new business demands.

2

Can your company afford sudden expenses upwards of USD 100,000? More and more firms are being hit by unexpected IT outages, and the costs they cause just keep rising. According to the annual Outage Analysis Report from the US Uptime Institute, 39 per cent of the companies surveyed in 2019 had paid more than USD 100,000 as the result of an outage. Just three years later, in 2022, the figure had risen to over 60 per cent. Of the affected companies, 15 per cent reported costs of over a million dollars. The chief causes of these outages were energy supply or network problems – not to mention the human factor. The institute based its study on a survey of 1,000 IT experts and data-centre operators.

### **Unforgiving digital business**

Software rules our world, so it's no surprise that the absolute number of IT outages is increasing. Software is not static. It is more like a living organism that is constantly adapting to its environment. It

faces new demands all the time, be they updates and maintenance work, changes in user behaviour, or new features from the competition. Yet a living organism reacts flexibly to change and can generally adapt well to new situations. Technology is not always so advanced in this respect.

Today, most companies operate at least a part of their business in the cloud. This ups the ante for IT systems in terms of scalability and reliability, all while they are becoming increasingly complex, and even distributed worldwide. Of course, online business is also changing customer expectations. It seems that firms have to be available 24/7 with a whole range of services if they want to remain competitive. That's why they are increasingly looking to Site Reliability Engineering (SRE) to achieve this, especially in the cloud. Essentially, it uses software engineering to largely automate system operation and make it fail-safe.

### **Faster, more secure, and more reliable**

In classic IT operations, incident management will handle the impacts of new releases. Once a problem is resolved, the incident is closed. In most cases this is manual work, using checklists. The problems arise where responsibilities are not clearly defined. SRE standardises and automates these tasks. Code and rules make the individual stages of work machine-readable. This replaces error-prone manual processes and also documents each step.

SRE has most to offer in cloud-native environments, where conventional IT administration processes soon reveal their limits. Cloud solutions are often technically complex, and changes to the infrastructure, platform, applications or services are errors and outages waiting to happen. Infrastructure as Code (IaC) offers a solution here. It involves specially written code replacing manual processes to attain a new level of automation. SRE in combination with the cloud makes companies both technically and organisationally flexible, creating the conditions for rapid growth. SRE means that new business requirements can be implemented faster.

### **Automation not manual intervention**

SRE creates a highly available system that can be operated with minimum manual intervention. Google once proclaimed that IT ops teams should be spending half their time on automation. At Ergon, it is more like 60 to 80 per cent thanks to the advent of better tools and services. Easily integrable monitoring and alerting services are one example here. If you already have these, it doesn't take much to expand their use to additional scenarios.

What does it make sense to automate? The answer depends on the system and the requirements. Repetitive tasks offer a good use case. Automated, rules-based monitoring can stop hard drives getting full, for example. It pre-emptively triggers data migration or deletes unwanted files in good time. The installation and rollback of new releases can also be handled automatically. Apart from routine work, it is also worth automating critical tasks, such as backups or file restores following data loss. Many companies do not actually perform daily backups, and the operations teams are not fully familiar with the steps involved. The Uptime Institute mentioned above points out that human error is generally the result of not implementing defined



### **Site Reliability Engineering in a nutshell**

SRE applies the methods and principles of software engineering to IT operations and IT infrastructure. The mission of dedicated SRE specialists or teams is to create extremely reliable, scalable software systems. SRE is often seen as a form of DevOps. In common with the DevOps approach, SRE closely links development and operations teams and ensures efficient development while maintaining operational stability. The aim of DevOps is to achieve quality and speed in development and delivery, while SRE focuses on a system's reliability for customers.

“Automated IT operations require software engineers.”  
\_Daniel Zeiter, Head of Technology, Ergon



“Site Reliability Engineering makes for lasting efficiency in cloud ops.”  
\_Laura Graf, Consultant, Ergon

procedures properly. Automation is the elegant way to rule this out.

SRE also improves forward-thinking, automated monitoring for services. With a proactive defence strategy, a company can deploy SRE to cut the risk of security incidents. SRE continuously monitors and analyses security data to identify and eliminate weak points before they can be exploited. Monitoring systems are a good example. SRE teams define what is to be monitored, and implement the relevant solutions in the code. Errors are then found and corrected at an early stage without the end user being any the wiser.

### Interdisciplinary teams

Software development generally involves a whole number of parties, with external partners often joining forces with internal ops and development teams. Regarding development and operations as isolated functions runs the risk of developing a silo mentality. SRE helps to find commonalities between the differing target scenarios: rapid new feature development on the one hand, and keeping software secure and stable on the other. With SRE, the ideal IT operation is managed by code that software engineers write during development, so you get Dev and Ops from a single source for efficient, fail-safe system management.

### Setting clear goals

A short outage will not necessarily determine the value of a system to its users. The higher the agreed service level, the longer the release time, and the higher the required budget. That's why 100 per cent availability is rarely the right goal. The important thing is the user's perspective. What they see as reliable can be identified on the basis of demand, supply-side competition, and benchmarks. Business and engineering teams together must work out which service level offers the best balance between business benefits and cost. The SRE team then translates this balance into measurable targets under service-level agreements. It also determines how these targets will be achieved, and takes swift action when things don't go to plan.

### SRE boosts business and IT

SRE has a positive effect on the whole company. Better IT operations reduce the error rate, and the time and cost of looking for bugs. If responsibility is bundled, there is no back-and-forth about whose job that is, making for fewer interruptions and outages, and faster response times. SRE has cut companies' IT outage times by 10 to 30 per cent, according to the Boston Consulting Group. Performance is 10 to 15 per cent higher, and software development two to five times faster than before. In other words, SRE makes companies more efficient. It is an investment that pays off. What's more, a good customer experience boosts long-term loyalty. Culturally, the approach encourages innovation and openness to change.

Gartner Consulting estimates that 75 per cent of companies around the world will be applying SRE principles by 2027. In fact, it should be part of strategy from the beginning, starting now, so that firms can build their IT teams accordingly. Although it isn't too late to automate IT ops even after moving to the cloud, it is generally more difficult culturally. SRE is designed to bring engineering and operations teams closer together and combine them where it makes sense. And that takes a great deal of cooperation, mutual trust and understanding.

With a shared understanding that operations are about software, too, reliability ceases to be onerous. Instead, it becomes a competitive advantage, and a key feature that every software needs nowadays. That is why companies should trust SRE as their autopilot, especially when heading to the cloud. />

**Interested  
in more?**

**Digitisation projects  
Change makers  
Tech trends**

**Order now**

[ergon.ch/smart-2023](http://ergon.ch/smart-2023)

