

WHAT IS...?

Site Reliability Engineering

www.itsmacademy.com



WHAT'S INSIDE:

What is Site Reliability Engineering?

Site Reliability Engineering Principles and Practices

SLOs, SLIs, and SLAs

Antifragility and Learning From Failure

SRE Doesn't Stand Alone

Make a Difference!

Want to Learn More?

RHI 2021 Salary Guide reports that an entry-level salary for a SRE starts at \$100,750, and an experienced professional in the role could earn \$199,000.

W

What is Site Reliability Engineering?

Site Reliability Engineering (SRE) is a discipline that incorporates aspects of software engineering and applies them to infrastructure and operations problems.

The concept, of site reliability engineering was first introduced at Google in 2003, when a team of software

"What happens when a software engineer is tasked with what used to be called operations."

Ben Treynor, Google

engineers was tasked with making Google's already large-scale sites more reliable, efficient, and scalable. The principles and practices they developed have continued to evolve and SRE is now increasingly being adopted by organizations running large-scale services. According to the DevOps Institute's [2021 Upskilling: Enterprise DevOps Skills Report](#), 22% of respondents have adopted SRE, up from 15% in 2020. This rise in SRE adoption shows the need to consider IT operations processes before services are released and while products are being planned, and not just as an afterthought; particularly in organizations that have adopted Agile and DevOps. DevOps is a set of practices, guidelines and culture designed to break down silos in IT development, operations, architecture, networking and security.

So, what's the difference between SRE and **DevOps**? It is often said that SRE is a narrower implementation of larger DevOps principles, with the primary distinction being that SRE's focus is on reliability in the production environment. Another distinction is that SRE is a job role.

Continue reading to learn more...

Email me the .pdf!



Site Reliability Engineering Principles and Practices

Like DevOps and other frameworks and disciplines such as Agile, **Lean**, and **IT service management (ITSM)**, SRE introduces new ways of thinking and working. This means culture change!

SRE principles and practices affect organizational culture.

The core principles and practices of SRE include:

- Accept that operations is a software problem - SRE uses software engineering approaches that focus on designing and building rather than operating and maintaining
- Manage service levels – in SRE, services are managed to service level objectives
- Reduce toil – if a task can be automated, it should be automated
- Leverage automation – don't automate a bad process – fix the process first
- Reduce the cost of failure – SREs view failure as an opportunity to learn
- Incentivize shared ownership – remove boundaries between dev and ops

The notion that individuals and organizations shouldn't automate a bad process is an important one. According to the DevOps Institute's 2021 Upskilling: Enterprise DevOps Skills Report, knowledge of process and frameworks is considered a key capability. This same study also found that 55% of respondents view process skills and knowledge as a 'must have' skill.

SLOs, SLIs, and SLAs

To manage services correctly, organizations must understand which behaviors really matter for that service, and how to measure and evaluate those behaviors. The concept of service level *agreements* is widely used in industry. Through the years, however, these 'agreements' have come to be viewed more as contracts with consequences, than an objective way to manage service levels.

SRE defines service level *indicators*, *objectives*, and *agreements*.

- **Service level indicator (SLI)** – a carefully defined quantitative measure of some aspect of the level of service that is provided
- **Service level objective (SLO)** – a target value or range of values for a service level that is measured by an SLI
- **Service level agreement (SLA)** – an explicit or implicit contract with users that includes consequences of meeting (or missing) the SLOs they contain

These measurements enable organizations to focus on metrics that matter, what values these metrics should have, and how the organization will react if it can't provide the expected levels of service.

What matters most in today's digital economy? Frequently used metrics associated with SLOs include availability, response time, and latency. These metrics should come as no surprise as they are all tightly related to the **customer experience**, which is a top priority for businesses.

To assist in influencing how the organization will react when service level objectives aren't being met, SRE introduces the concept of error budgets. An **error budget** is a clear, objective metric that determines how *unreliable* a service is allowed to be within a specific time period. An error budget is typically expressed as 100 percent minus the SLO of the service. For example, a 99.9 percent SLO has a 0.1 percent error budget. Every month, this allows for 43 minutes of outages. New feature releases, patches, planned and unplanned downtime need to fit into these 43 minutes.

SRE practices encourage organizations to strategically burn the budget to zero every month, whether it's for feature launches or architectural changes. This way teams know that they are running as fast as they can (**velocity**) without compromising availability. Having said that, failure to hit an SLO must have consequences.

ERROR BUDGETS

A tool that SRE uses to balance service reliability with the pace of innovation.

Error budget policies spell out the activity a team takes when they've exhausted their error budget for a particular service in a particular time period (e.g., no new features launches, prioritizing remediation work).

Antifragility and Learning From Failure

"There is no such thing as failure. There are only results. It's time to stop beating yourself up and start realizing that everything you do is a success or a learning experience."

Tony Robbins, author & life coach

Many organizations view failure as a bad thing. Something to be avoided at all costs and something that involves repercussions for those involved. The end result is a culture that fears failure. These same organizations often advocate big bang, large-risk deployments, proceeded by lots of overly rigorous processes, all in an effort

to eliminate risk and uncertainty. The very kind of deployments that often cause outages and have – wait for it – lots of difficult to diagnose failures.

The 5 whys is an iterative interrogative technique used to explore the cause-and-effect relationships underlying a particular problem. Let's apply this technique to the problem of how to prevent outages.

- We need to be better at breaking things – Why?
- We need to understand how things work together – Why?
- So when things break, we know how to fix them – Why?
- So we can prevent things from breaking – Why?
- To reduce downtime – Why?
- To ensure we make money*

* Replace money with whatever value means in your organization.

Embracing failure will help improve key antifragility metrics such as mean time to detect (MTTD) incidents, mean time to recover (MTTR) components, and mean time to restore (MTRS) services.

Antifragility is a property of systems that increases its capability to thrive as a result of stressors, shocks, volatility, noise, mistakes, faults, attacks, or failures. While counterintuitive, introducing failure makes it possible to optimize monitoring capabilities, to determine where automation can be used to auto-recover, and to determine where more resilience is needed to shore up single points of failure.

So how do we shift from a culture that fears failure to one that embraces and learns from failure?

- Enable “The Third Way” – The Third Way of DevOps encourages organizations to foster a culture that embraces continual experimentation, taking risks and learning from failure.
- Benchmark against the “Westrum Model” – The response to failure shapes the culture of an organization. If failures lead you to ask questions about what caused the failures in a blameless way (as is the case with a generative, performance-oriented culture) and how you can keep them from happening again in the future, you've improved not only your technical systems, but also the skills of your teams, your processes, and your culture.
- Introduce “Fire Drills” – typically an audit requirement, fire drills build on the concepts of business continuity planning (BCP) and disaster recovery (DR).
- Introduce “Chaos engineering” – introduce resilience testing and pinpoint areas where more resilience may be required.

Chaos engineering is the discipline of experimenting on a software system in production in order to build confidence in the system's capability to withstand turbulent and unexpected conditions. Chaos engineering enables organizations to:

- Assess and improve monitoring and logging capabilities
- Identify dependencies
- Improve error handling and recovery (manual-to-automated)
- Learn from “real” failures

The organization, its employees, and its customers benefit as organizations are able to:

- Minimize the impact of failures on the customer experience (and hence revenue)
- Introduce automated self-healing capabilities resulting in fewer “on call” incidents

Why Organizations Embrace SRE

According to Gartner, the average cost of service downtime is \$5,600 per minute. Because there are so many differences in how businesses operate, downtime, at the low end, can be as much as \$140,000 per hour, \$300,000 per hour on average, and as much as \$540,000 per hour at the higher end.

Combine this with the fact that as more and more people move online and become reliant on services, the impact of an outage increases.

Downtime may be commercial (e.g., lost sales), but can also impact lives (e.g., if a service is serving as an information system for first responders, or delivering benefits to claimants). There may also be financial penalties for being offline (e.g., in the case of financial institutions).

GOING VIRAL

Social media can make outages “go viral” which can have a massive impact on brand reputation and ultimately business performance.

These reasons may seem compelling enough but it is important to keep in mind that SRE, like many other modern frameworks and techniques, is going through the “hype curve”. Beware the hype!

So, what organizations should embrace SRE? Any organization that is witnessing growth can benefit from SRE principles and practices.

- Platform growth (large volumes of users, irregular data flows, legacy-to-modern architectures)
- Scope growth (new products and services)
- Ticket growth (volume of incidents, outages, requests, toil)

Toil is a kind of work tied to running a production service that tends to be manual, repetitive, automatable, tactical, and devoid of enduring value.

SRE Doesn't Stand Alone

Organizations, particularly digital organizations, are constantly looking at their end-to-end value streams and looking for ways to optimize their performance. Sometimes the plethora of seemingly competing frameworks can look like a maze. In reality, SRE can and often does interact with other frameworks, and can be highly complementary.

- Agile – SRE teams can operate in an Agile way – using frameworks like Scrum and Kanban
- DevOps – SRE helps to further breakdown organizational silos and extends the reach of the DevOps pipeline
- Lean – SRE brings the ‘wisdom of production’ to help improve flow across the IT value stream
- ITSM/ITIL – Like ITSM, SRE aims to **‘optimize and automate’** and emphasizes continual improvement and learning from failure

SRE provides businesses a wide focus on stability and reliability. The engineering solutions introduced by SRE help organizations to scale services without the need to scale the organization.

Make a Difference!

Any initiative that introduces new ways of thinking and working will affect organizational culture. Effective communication plans, training, guiding principles, and clear policies and procedures are all needed to achieve the desired performance outcomes and enable collaboration between the many different people involved in continual improvement.

Contribute to your organization’s SRE adoption effort by expanding your knowledge of best practices and by enthusiastically using what you learn to lead implementation and improvement activities.

BE A CHANGE CHAMPION!

Culture change and progress cannot happen without the support of people like you. **Take action!**

Want to Learn More?

Training helps individuals and organizations build and maintain their capabilities. Training also provides individuals the knowledge, skills and information needed to fill their role in the organization or achieve their career goals, along with a place to test and develop the confidence to use these skills in the workplace.

ITSM Academy's [DevOps Campus](#) provides the courses you need to build a solid foundation and sharpen your skills as a DevOps practitioner.



According to the [2021 Upskilling: Enterprise DevOps Skills Report](#), 47% of respondents voted SRE skills as a must-have skill in the process and framework skill domain, compared to 28% in the 2020 year's survey.

Our [Site Reliability Engineering Foundation](#) course provides an introduction to the principles & practices that enable an organization to reliably and economically scale critical services. The course introduces a range of practices, methods, and tools for improving service reliability through a mixture of automation, working methods and organizational re-alignment.

[Click here to learn more about Site Reliability Engineering Foundation.](#)

[Visit our website for our full class catalog and public class dates.](#)

Contact us to schedule time with a subject matter expert.

+1-954-491-3442

<http://itsmacademy.com>

info@itsmacademy.com

Additional Resources:

- [ITSM Professor Blog](#) - a WEALTH of knowledge published weekly since 2008
- [Webinar Archives](#) - Monthly since 2007
- [ITSM Academy Resource Center](#)



ITSM Academy

We are a female owned small business, established in 2004. Our extensive catalog contains accredited and sustainable IT Service Management (ITSM) education and advice including; ITIL®, DevOps, Process Design (CPDE), Agile, Site Reliability Engineering (SRE), Value Stream Mapping (VSM) and Experience Level Agreement (XLA). Our business values are founded on trust, loyalty, professionalism and long term relationships.

...educate and inspire is not just our corporate slogan, it speaks to our core mission and goal.



Follow our founder and CXO, Lisa Schwartz, on [LinkedIn](#).

Instructors

Every ITSM Academy instructor is certified to the highest levels in the areas they train. They have years of hands-on IT practitioner experience, enabling them to effectively intertwine theory and real-life stories and scenarios. Using the highest quality content, this engaging training style encourages active group participation, allowing all learners to bring from class a wealth of practical and actionable knowledge.

Accreditations

All of ITSM Academy's certification courseware is developed or enhanced in-house and is accredited by independent, international organizations where applicable.



Game On! - Interactive Learning

Involves students in active learning, using the engaging qualities of a game, fueled by our subject matter experts.

Courseware Licensing (all developed or enhanced in house)

In addition to our public and corporate/onsite training, our courseware is available for licensing / co-branding under our flexible licensing program, including Train-the-Trainer (for qualifying organizations).

my.itsmacademy.com (digital portal)

Extends the learning experience with games, videos, exercises, sample exams, and course materials. It also provides instructors a vast repository of information and guidance to successfully prepare for and teach our courses.

Professional Education Hours (CPDs/PDUs/CPEs/CEUs):

ITSM Academy is proud to make it possible for individuals who attend our classes to earn professional education hours. (e.g., CPDs, PDUs, CPEs, CEUs). These professional education hours can be submitted to associations such as PeopleCert, the Project Management Institute and ISACA, if applicable.



[Read More](#)

The Story of the Academy

Today, ITSM Academy is widely recognized for its expertise in multiple IT frameworks (ITSM, ITIL, Process Engineering (CPDE), DevOps, Agile Service Management, Lean) and, more importantly, how they work together. But that's not where we started.

+1-954-491-3442

<http://itsmacademy.com>