NetImpact Strategies

Meaningful Measures of Operations and Service Delivery

How SRE Transforms Metrics into Powerful Tools for Operational Excellence

7600 Leesburg Pike, West Building, Suite 140, Falls Church, VA 22043
(703) 559-3280 | www.netimpactstrategies.com

INTRODUCTION

As agencies navigate the complex terrain of delivering critical services and safeguarding sensitive data, they often find themselves embarking on a relentless quest for ways to build an infrastructure foundation with rock-solid reliability and streamlined operations. Reliable IT operations are required for the success of any organization. By ensuring reliability in systems, organizations can improve business continuity, enhance employee productivity, improve customer satisfaction, reduce costs, enhance security, and simplify compliance.

A powerful and transformative approach has been quietly reshaping the commercial tech world and is becoming a game-changer for federal IT: Site Reliability Engineering (SRE). SRE is a pragmatic, results-driven approach that can catapult your IT operations to the next level and introduce revolutionary possibilities in elevating the reliability and performance of the organization's digital services while simultaneously fostering a culture of proactive problem-solving and innovation.

Google Cloud's 2022 State of DevOps Report states that organizations with mature SRE practice deploy code 20 times more frequently than those without an SRE practice, have a change failure rate that is 10 times lower, have a mean time to recovery that is 100 times faster, along with deployment success rate that is 99.9%.

Site Reliability Engineering (SRE) Overview

Site Reliability Engineering (SRE) is a discipline that blends aspects of software engineering with operations to create a highly reliable, scalable, and efficient IT environment. SRE is built on a foundation of automating tasks, real-time monitoring and alerts, improving incident responses, optimizing operations, building reliable systems, and enhancing collaboration within teams. Unlike conventional IT operations that focuses on reactive firefighting, SRE promotes a proactive, preventive approach that aims to eliminate recurring issues, minimize downtime, and optimize system performance. SRE is intrinsically tied to the principles of DevOps because it envisions a seamless blend of software engineering and operations expertise. Instead of traditional IT operations teams handling tasks manually, SRE teams empower engineers and operations personnel to leverage cutting-edge tools and automation, propelling problem-solving and system management at new speeds.

In essence, the major distinguishing aspect of SRE from conventional IT operations is the “what” and “how” of automation. SRE teams assume responsibilities that were once the domain of IT operations and streamlines those duties with code-driven solutions. Automation becomes the cornerstone of this transformative methodology, enabling SRE teams to tackle challenges swiftly and maintain the reliability of production systems. By shifting from manual tasks to a scalable, software-driven approach, SRE not only facilitates the creation of highly reliable software systems but also empowers organizations to manage massive infrastructures with greater ease and sustainability. SRE practice’s broad responsibilities are to optimize and elevate system availability, latency, performance, efficiency, change management responsiveness, monitoring, emergency response, and capacity planning of services.

SRE's Role in Measuring and Improving Reliability in Systems

Customer experience and satisfaction are paramount when measuring success of any organization. To achieve success, setting up clear benchmark, accountability, continuous improvement, and reliability in systems are crucial. SRE team drives efficiency in a system through meaningful measures for operations and maintenance activities. It uses three key concepts to measure and improve the reliability of systems:

1. Service-Level Indicators (SLIs)
2. Service-Level Objectives (SLOs)
3. Service-Level Agreements (SLAs)

SLIs are the key measurements of the system's availability, SLOs are the set goals for how much availability is expected out of a system, and SLAs explains what happens if a system doesn't meet its SLO. These key concepts and the measurements' relationship to each other sets clear performance targets, aligned with user expectations.

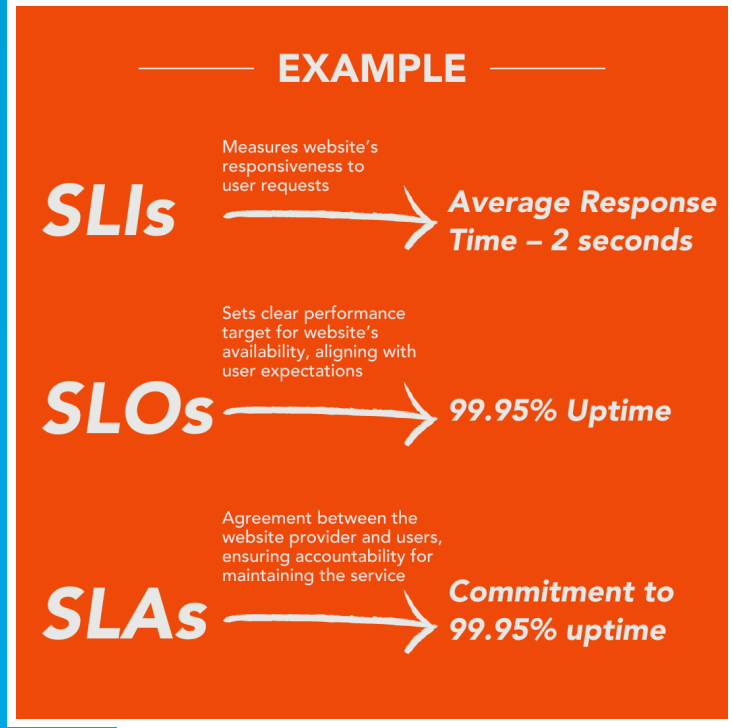

SRE focuses on building intrinsic performance and reliability into IT systems, often resulting in exponential performance increases in key measurements.

Some of these key measurements include:

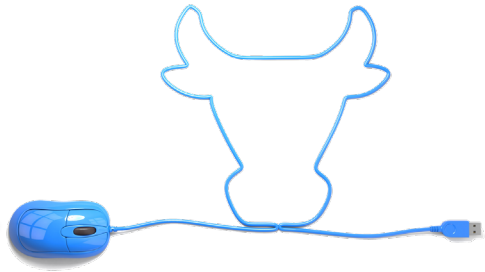

Change Velocity

What volume of change can be successfully introduced into the systems in a given period of time?

Request Error Rate

What percentage of requests result in implicit or explicit errors, within a given period?

Mean Time To Failure (MTTF)

How long is a system expected to run before failure?

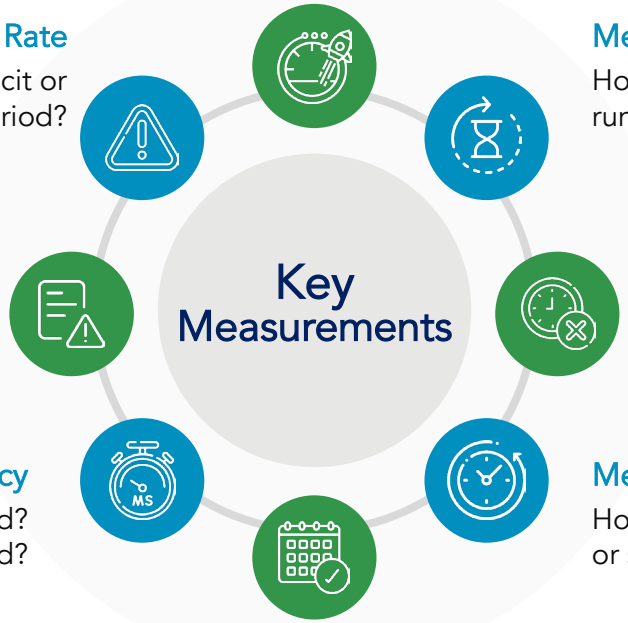

Request Throughput

How many error-free requests can a system process within a given period?

Mean Time Between Failures (MTBF)

How much operational time exists between system failures?

Request Latency

How long does it take a successful request to be served?
How long does it take for a failed request to be served?

Mean Time To Restoration (MTTR)

How quickly is service restored following an outage or service degradation?

System Availability

What percentage of time is the system available and functioning during the times users expect it to be available and functioning?

Enhance Customer Experience by Employing SLIs

A Service-Level Indicators (SLI) is a vital component of Service Level Management, providing precise measurements to assess the performance of a service. SLIs are quantitative metrics tracking specific aspects of service delivery, enabling organizations to monitor and improve service quality by providing us clues into how well our IT systems, processes, and personnel are performing.

For instance, the United States Department of Agriculture (USDA) oversees laws and programs related to farming, forestry, rural economic development, and food. Their agencies are also responsible for keeping America's farmers and ranchers in business. Let's say they are deploying an application on their website to enable ranchers to submit funding request for livestock loss due to the extensive California wildfires. USDA needs to collect information and process emergency relief payment as a part of the *Extending Government Funding and Delivering Emergency Assistance Act (P.L. 117-43)* and in accordance with the disaster relief program requirements, like application processing timeframes.

The Department's SLI would be defined with the user experience and expectations of the ranchers affected by these type of disasters in mind. USDA would be interested in metrics related to the latency of the web application so they may be using Time to Interactive (TTI) measures to understand how long it takes for the page to fully load or render until it's ready for user input – gaining insight into application responsiveness.

Federal government agencies often utilize SLIs to evaluate the reliability and effectiveness of their digital services. As per the United States Digital Service Playbook, metrics like response time, system availability, and incident resolution rate – e.g.: achieving 99% within 4 hours – serve as key SLIs to monitor service performance. By employing SLIs, federal agencies can proactively address service bottlenecks, enhance user experiences, and uphold their commitment to delivering seamless and efficient services to the public.

However, implementation of SLIs can be a challenging task for federal agencies operating legacy application portfolios, built on top of aged infrastructure and manual tools. There are a few pre-requisite components you'll want to have in place to ensure you can properly collect and attribute SLIs.

In essence, SLIs measure an aspect of what the end user or customer is experiencing.

Common Service Data Model (CSDM)

A Common Service Data Model provides a common language of entities and relationships that brings disparate IT teams together. An organizations' Enterprise Architects (EA) are typically charged with developing this data model, often referred to as an EA meta-model. A popular CSDM is the meta-model prescribed in The Open Group Architecture Framework (TOGAF). While they make an incredible foundation to begin from, rarely does a framework's prescribed metamodel match an organization's needs completely so significant customization is often necessary. Adoption of CSDM's often require changing deeply embedded knowledge and biases. Tooling can often serve as lever to efficiently influence these changes. If an organization leverages ServiceNow, the ServiceNow CSDM is an excellent option for accelerating the adoption of a common services data model because the data model is engineered into all ServiceNow's IT product offerings.

Configuration Management Database (CMDB)

To measure the performance of any system, you must know the parts of which it is comprised. IT systems are comprised of layers of complex inter-dependent components. It's an impossible job to keep up on these ever-changing inventories of components and relationships. A modern Configuration Management Database (CMDB) solves this problem. By leveraging a pre-built data model and automated network scanning tools, a CMDB discovers and catalogues all the bits and bolts of your IT environment. These underlying bits and bolts are the very things we must implement our SLI measures against.

Monitoring and Observability Infrastructure

IT systems are chatty little things – as they should be. They want to tell everyone about everything they have going on but to get any value from this behavior, somebody, or something must be there to listen. The information necessary to measure SLIs is present within this data. Monitoring and Observability tools collect all the communication our systems are storing locally and aggregate them to a central repository where those messages can be correlated, enriched, analyzed, and actioned upon.

Exceed Performance Benchmarks by Establishing Meaningful SLOs

Service-Level Objectives (SLOs) are the desired targets or goals set for SLIs. They are objectives we set for clear performance targets and how much availability we expect out of a system. These are set based on user expectations and aligns with business goals. By continuously monitoring service performance against SLOs, organizations gain insights into whether their services are meeting user needs and if adjustments are required. Defining and communicating SLOs to users plays a crucial role in managing expectations for service performance. By transparently establishing these objectives, service owners can mitigate unwarranted complaints from users, such as concerns about slow service performance. Without clear and explicit SLOs, users might develop subjective expectations regarding desired performance, which may not align with the intentions of the service designers and operators. By setting specific SLOs, organizations can bridge this gap in expectations and foster a common understanding of performance targets, ultimately enhancing user satisfaction and aligning user expectations with the service's capabilities.

Following our USDA example, the Department may set a responsiveness SLO around ensuring that the loan request portal can quickly assess whether the rancher is eligible for relief payment. For this part of the web-based application, ranchers must estimate their loss declaration and respond to other criteria questions. USDA's SLO could be that the portal returns the eligibility results in 90 milliseconds for 95% of all requests. The various SLI measures for latency captured around TTI, load time, and page speed would demonstrate whether the portal behavior is meeting the objective. Another SLO the Department may focus on may be from standards already set, like maintaining 99.95% system availability for essential government applications.

SLO metrics provide tangible benchmarks for measuring success and optimizing resource allocation.

Ensure Accountability by Leveraging SLAs

Service-Level Agreements (SLAs) are formal agreements between service providers and consumers that specify the expected level of service performance. It is a crucial component of effective service management, outlining the agreed-upon performance expectations between a service provider and its customers. SLAs set clear benchmarks for service quality, availability, and responsiveness, ensuring that both parties share a mutual understanding of what to expect. If these commitments are not met, the SLA could define penalties or compensations.

The Department may hold an SLA with its cloud service provider to guarantee 99.9% uptime for critical systems. This means the service provider commits to ensuring that USDA's systems are available and accessible to ranchers attempting to access the portal 99.9% of the time. Another example could be an SLA with the vendor overseeing this system to respond to critical incidents within 1 hour as the vendor's commitment to timely incident resolution, which helps ensure ranchers have reliable assistance from USDA.

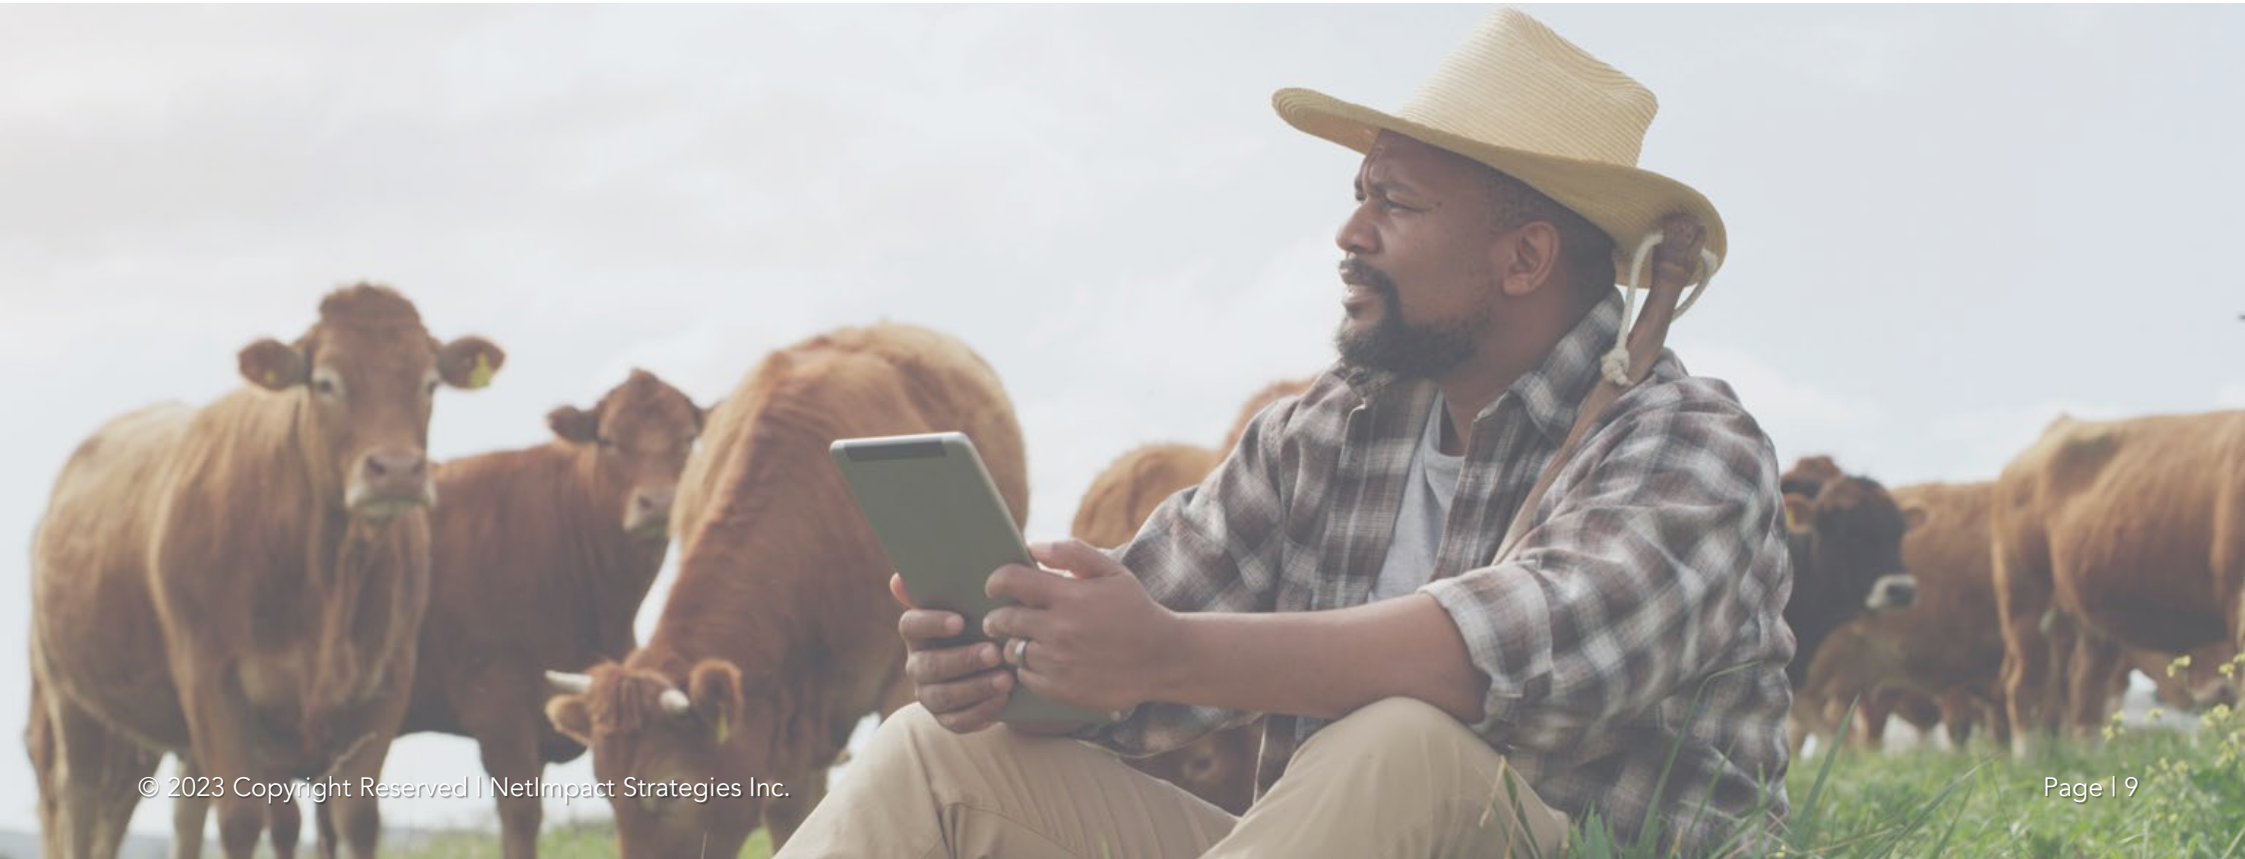

What makes SRE different?

SLIs, SLOs, and SLAs have been a cornerstone of operations and maintenance practices for a considerable time – what makes them different within the context of SRE? SRE's fresh perspective comes from new dimensions it applies to these measurements, which transforms their roles from being mere benchmarks to dynamic tools for proactive strategies and unparalleled reliability. This transformation acts as a driver and encourages streamlined operational and maintenance activities across the organization.

Here are the key differences that set SRE apart from traditional approaches:

Proactive vs. Reactive

Contrary to the traditional approach of reacting to incidents in a predefined timeframe, SRE emphasizes a proactive approach. It starts with defining SLIs and associated SLOs, which are ambitious yet achievable targets for reliability and performance. SRE teams work tirelessly to ensure systems meet these objectives, preventing issues before they impact end-users.

Error Budgets

An error budget quantifies the allowable amount of downtime or errors a service can experience while still fulfilling its SLOs. This shift in thinking recognizes that perfection is unattainable, and some issues will inevitably arise. Error budgets provide a mechanism to strike a balance between rapid innovation and system reliability. If the error budget is depleted, new changes may be temporarily halted until reliability improves.

Data-Driven Decision Making

SRE relies heavily on data analysis and continuous monitoring to measure system performance against SLOs. This data-driven approach allows teams to make informed decisions, identify patterns, and proactively address potential vulnerabilities, leading to more resilient systems.

Blameless Postmortems

In traditional SLA-driven environments, the focus on meeting targets can sometimes lead to a culture of blame when incidents occur. SRE champions a blameless postmortem culture, where teams analyze incidents in a non-judgmental environment to learn from mistakes and improve systems collaboratively. The focus is on continuously improvement as a team.

By maintaining a data-driven and continuous monitoring approach, SRE teams make informed decisions, identify patterns, and proactively address potential vulnerabilities, leading to more resilient systems, and foster a collaborative and blame-free environment. This shift in approach elevates SLAs from mere contractual obligations to powerful tools for driving reliability and operational excellence. It also re-contextualizes the technology around delivering service to the customers who are relying on the tools for their work and, in many times, their livelihood.

For federal agencies navigating the complexities of IT infrastructure, SRE emerges as the beacon of possibility. With a data-driven focus on system performance, agencies can pinpoint areas for enhancement and allocate resources with precision. The metrics gathered through SRE practices offer an unparalleled opportunity for agencies to elevate their IT landscape, achieve robust system availability, and mitigate critical incidents with efficiency and confidence. As the landscape of federal IT operations evolves, the insights harnessed from SRE hold the potential to usher in an era of enhanced reliability and seamless service delivery.

About NetImpact Strategies Inc.

Digital transformation is essential for the federal government to fulfil their missions efficiently and effectively. However, complex regulations, legacy systems, and a diverse workforce pose unique challenges.

NetImpact is a NextGen digital transformation leader disrupting how technology is applied to deliver mission value. Our DX360°® framework provides a tailored, cohesive, and holistic approach to digital transformation that addresses the specific needs of the federal government. Encompassing Strategy, Culture, Technology, and Data, DX360°® enables the federal government to successfully navigate digital transformation complexities and deliver high-quality digital mission enablement.

Our team understands the challenges of managing and modernizing your agency. Building on our decades of experience in developing technology that integrates seamlessly into your ecosystem, we developed a portfolio of powerful, yet easy-to-use DX360°® apps that are designed to help you tackle your mission's most pressing needs and comprehensive services to guide the journey. Our DX360°® suite of high-performance digital solutions drives impact by delivering agile, outcome-focused results to securely transform client operations and accelerate mission outcomes at a fraction of the cost of traditional projects.

Contacts

NetImpact Strategies HQ

7600 Leesburg Pike
West Building, STE 140
Falls Church, VA 22043
P: (703) 559-3280
info@netimpactstrategies.com

Follow Us

NetImpact Strategies