

DELIVER ON SERVICE RELIABILITY WITH BMC HELIX AND SRE



Site reliability engineering (SRE) is a set of principles and disciplines that helps to achieve reliability for the services a company provides. This article offers an introduction to SRE principles and explains how [BMC Helix](#) can help maximize and optimize service and operations reliability.

SRE: “The consummate DevOps how-to-manual”

Site reliability engineering (SRE) is a set of principles and practices that incorporates aspects of software engineering and applies them to IT infrastructure and operations (I&O) problems. The main objective of SRE is to create ultra-scalable and highly reliable software systems. A site reliability engineer spends approximately:

- Fifty percent (50%) of their time on operational activities such as incident management, on-call support, and manual intervention
- Fifty percent (50%) on development tasks like feature development, auto scaling, and manual task automation

The SRE discipline helps break down silos between software engineering and operations activities within the organization.

(Read our [full SRE explainer](#) & [comparison to DevOps](#).)

There have been multiple attempts at defining a canonical list of SRE principles, but these have often lacked consensus. The following characteristics, however, are usually included in most definitions:

- Automating or eliminating of repetitive tasks cost effectively

- Pursuing only the degree of reliability considered business-critical, which implies established understanding of that metric
- Designing systems with a bias towards reducing availability, performance, and latency issues
- Increasing visualization capabilities and a wide variety of questions about the system can be asked at any time

SRE is generally characterized by seven principles, including:

1. Embracing risk
2. Developing service-level objectives
3. Eliminating overhead
4. Monitoring
5. Release engineering
6. Automation
7. Simplicity

These principles focus on key metrics of mean time to failure (MTTF) and mean time to repair (MTTR). Organizations need to implement solutions that not only detect problems proactively but also can scale across the entire enterprise and intelligently automate for consistent service performance and reliability. So how does BMC Helix help organizations achieve SRE excellence?


Connecting across domains for visibility, observability, and intelligent actionability: BMC Helix Platform

Enterprise organizations are running and reinventing themselves as they become an [Autonomous Digital Enterprise](#) (ADE). They are striving to achieve these three key traits of success:

- Agility
- Customer centricity
- Actionable insights

Organizations turn to BMC Helix to help them reach these success goals for their service and operations needs. BMC Helix solutions are powered by the BMC Helix Platform, which provides an open, scalable, and unified foundation for increased organizational efficiency, productivity, and innovation.

BMC Helix was architected to take full advantage of the requirements needed for SRE as well as seamless integration across the entire enterprise.



BMC Helix is optimized for

SRE-centric capabilities from the platform level to help companies manage services more efficiently and at higher quality

Principle 1: Embracing risk

Embracing risk begins with identifying the reliability acceptable by customers. No service or product can ever be 100 percent reliable, and while customers generally accept this fact, they do have a maximum tolerance level. At the same time, improving reliability always incurs cost.

Embracing risk can help weigh the cost versus risk for site reliability engineers and customers. Site reliability engineering aims to strike a balance between the risk of unavailability and the goals of rapid innovation and efficient service operations so that users' overall happiness—with features, service, and performance—is maximized.

BMC Helix delivers key capabilities:

- Intelligently understand risk by providing a holistic view and analysis of the process of monitoring, analyzing, and extracting insights from the SRE ecosystem.
- Implement this SRE principle by providing continuous visibility on customer service-level indicators (SLIs), service-level objectives (SLOs), and error budget.
- Automatically track availability of application instances and services, and compare actual availability against SLOs.
- Dynamically set the risk level of a change request based on team performance (outages), resulting in increased control if outages exceed a certain limit.

Principle 2: Setting service-level objectives

SLOs essentially translate customer satisfaction into internal goals. An SLO is composed of an SLI, a duration, and a target. For example, an SLI might be the ratio of the number of responses with HTTP

code 200 to the total number of responses. The duration is the total time in which a metric is measured. This period can be calendar-based (for example, from the first day of one month to the first day of the next) or a rolling window (for example, the previous 30 days).

The target might be the desired percentage of good events to total events (such as 99.9 percent) that you expect to meet for a given duration.

Here are some SLO examples:

 TYPE OF MEASURE	EXAMPLE SLO REQUIREMENT	MEASUREMENT PERIOD
Availability	Application will be available 99.95% of the time	Over a year
Service desk response	75% of help desk calls will be answered in less than a minute 85% of help desk calls will be answered within 2 minutes	Over a month
Incident response time	99% of severity 1 tickets will be resolved within 3 hours 98% of severity 2 tickets will be resolved within 8 hours	Over a quarter
Response time	85% of transfer control protocol (TCP) replies within 1.5 seconds of receiving a request 99.5% of TCP replies within 4 seconds of receiving a request	Over a month

Setting these goals requires intuition, experience, and an understanding of what users expect so that organizations can set achievable SLIs, SLOs, and SLAs.

These measurements describe the basic metrics that matter and how to measure them, what values we want those metrics to have, and how we'll react if we can't provide the expected service. To set the SLOs effectively, you need to be able to understand customer pain points and budget accordingly. You may also have to modify them as customer experience dictates.

BMC Helix empowers SRE engineers:

- Manage complex infrastructures including [multi-cloud](#), hybrid cloud, and distributed systems, and use [artificial intelligence for service \(AISM\) and operations](#) (AIOps) to predict and proactively address issues before they impact service.
- Leverage a unified view of operations data (metrics, events, and topology) from disparate sources and consolidate them to create AI-driven actionable insights.
- Build SRE dashboards that provide amalgamated visibility into SLIs and SLOs to meet service quality and availability needs.
- Implement extensive monitoring coverage, ensuring that all the SLIs are monitorable and keeping SLOs and error budget updated in real-time.
- Use an intuitive dashboard for a bird's-eye view into the SRE environment, which helps maintain the reliability promised to customers.

Principle 3: Eliminating toil

Eliminating toil means reducing the number of mundane tasks performed by a site reliability engineer. With increasing automation, many of the activities performed by an SRE this year won't need to be accomplished manually next year. Key aspects of eliminating toil are to identify the repetitive tasks and automate those activities. The benefit of this is that organizations free up resources and upskill or simply better employ them.

BMC Helix gives insights:

- A policy-driven automation broker that enables SRE teams to quickly identify valuable automation opportunities and to implement hyperautomation for mundane tasks.
- Identification of requests that are addressed manually or knowledge searches that are performed without success.
- AI-based identification of recurring incidents that consume the most manual work for the [service desk](#), transitioning those to the problem investigation process so that the [root cause](#) can be identified and resolved.
- AI-led, automated recommendations for problem fixes that can reduce agent searches for [knowledge articles](#) and similar tickets.
- A wide variety of options to automate human-intensive operational tasks by integrating with external automation tools of customer choice.
- Continuous detection of the state of the infrastructure and [artificial intelligence/machine learning](#) (AI/ML) algorithms and policies that generate recommended automated tasks that SRE engineers can implement.

Principle 4: Monitoring

[Monitoring](#) is foundational to SRE principles. In an SRE world, organizations need to monitor, manage, and optimize the end-to-end performance and availability of infrastructure and applications across increasingly complex and hybrid IT environments. At the same time, they need to support the agility, speed, and scalability required by SRE initiatives.

In "[Monitoring Distributed Systems](#)," the author states that metrics and structured logging are the two data sources best suited for SRE monitoring needs as they can provide insights for trend analysis and troubleshooting, which point to the root cause.

"The Four Golden Signals" are the most common metrics used to measure site reliability:

- Latency: the time it takes for a service to respond to a request
- Traffic: the amount of load a service is experiencing
- Error rate: how often requests to the service fail
- Saturation: how much longer the service's resources will last

The BMC Helix Platform meets all the requirements for a modern monitoring strategy and provides a cohesive platform to deliver all the phases of AIOps like observe (monitoring), engage (ITSM), and act (automation).

With BMC Helix, blindspots are eliminated:

- Implement real-time visibility into hardware, software, services, and dependencies across multi-cloud and hybrid-cloud environments.
- Leverage discovery capabilities designed to handle the complex management of hyper-converged infrastructures, software-defined storage and networks, containers, and cloud services.
- Reduce the [mean time to resolution](#) (MTTR) and improve performance and availability using the solution's predictive, AI-driven monitoring, event consolidation, and correlation.
- Predict and proactively address potential issues across complex IT infrastructures using advanced analytics and machine learning to detect anomalies and identify root causes.
- Consolidate data from disparate sources with BMC Helix's open API integration capability.
- Leverage anomaly detection (univariate or multivariate) that proactively triggers events and notifications based on one or more metrics behaving abnormally. Proactive notification of issues enables corrective action before there is an impact to service.
- Perform root-cause analysis, [postmortem analysis](#), and service management using built-in AI and ML capabilities. BMC Helix also predicts and prevents outages while reducing event noise and driving significant improvement in MTTR.
- Enable dynamic service modeling features that generate a topology view of services leveraged for service-centric monitoring. Probable cause analysis provides rapid root-cause identification, which will significantly improve the mean time to acknowledge (MTTA) and reduce the MTTR by creating actionable insights and preventing alert fatigue.
- Access a 360-degree view into SRE teams, offering real-time insights across all business and IT services.
- Handle the volume, variety, and velocity of rapidly proliferating machine data so engineers can search, analyze, and visualize while also applying AI/ML capabilities to create operational insights for the SRE team.
- Take advantage of post-mortem review capabilities in the incident and problem management process. Blameless post-mortem reviews can be conducted with results being documented and shared with other organizations.

Principle 5: Automation

Automation frees teams from repetitive tasks, such as scaling resources as the workload demands and optimizing resources by analyzing workload needs.

BMC Helix delivers intelligent automation:

- Acts as a bridge between all third-party automation solutions or any custom scripts with a vendor-agnostic automation broker.
- Augments automation tasks, such as spinning up servers and configuring systems, etc., during the deployment phase.
- Auto-scales resources for workload demands as per the business needs without human intervention.
- Enables SRE teams to supply the technology resources required to deliver modern applications and service assurance. No matter how dynamic, complex, or diverse the IT environment is, the SRE team can use the insights enabled by AI and ML to optimize the

operational expenditure (OPEX) significantly.

- Provides at-a-glance insight into the technology resource usage for business services so they can analyze capacity at the service pool or deployment level.

Principle 6: Release engineering

Release engineering looks at building and [deploying software](#) in a consistent, stable, and repeatable way. Running reliable services requires reliable release processes. Site reliability engineers need to know that the binaries and configurations they use are built in a reproducible, automated way so that releases are repeatable and aren't "unique snowflakes." Changes to any aspect of the release process should be intentional rather than accidental.

BMC Helix removes the “guess work”:

- Reduce the complexities of customer support, change, and release management, as well as asset management processes and activities, by making them simple, integrated, and efficient.
- Enable agile DevOps organizations to maximize the delivery and overall quality of service while ensuring governance and compliance.
- Achieve comprehensive change and release management with minimal risks and further the goal of release lifecycle automation, which defines how release requests are initiated and managed through fulfillment by coordinating change implementations.
- Leverage automated and contextual collision detection and impact analysis, as well as enhanced risk analytics that automate routine changes that deliver critical information needed for agent decisions.

Principle 7: Simplicity

Simplicity is an important goal for site reliability engineers, as it strongly correlates with reliability. Simple software breaks less often and is easier and faster to fix when it does break.

Simplicity for a site reliability engineer is a holistic and end-to-end approach to reliability. It should extend beyond the code itself to the system architecture, tools, and processes used to manage the software lifecycle. The SRE team is in an excellent position to identify, prevent, and fix sources of complexity, whether they are found in software design, system architecture, configuration, deployment processes, or elsewhere, because of their end-to-end understanding of the systems.

BMC Helix delivers a single-pane-of-glass for the entire enterprise:

- Offers a 360-degree view of real-time insights across all business and IT services.
- Enables comprehensive solutions to achieve the simplicity principle of the SRE, even if the organization's software or infrastructure is complex.
- Covers all the aspects of SRE by providing corrective, preventive, condition-based, and predictive maintenances.

BMC Helix is built on a unified platform that automates and enables

the work of service and operations teams

To summarize, **BMC Helix** is the only unified solution with modern, containerized architecture and a deployment model “of choice” available in the market today. These capabilities enable organizations to adopt and implement SRE principles seamlessly.

BMC Helix provides key functionality and differentiation with its ServiceOps approach, which enables collaboration, communication, and orchestration across the enterprise for faster, higher quality service. BMC continues to deliver innovation on SRE-centric capabilities to help organizations align with SRE principles for managing their services more efficiently.

Related reading

- [BMC Service Management Blog](#)
- [BMC IT Operations Blog](#)
- [BMC AIOps Blog](#)
- [SRE vs ITOps: Are SRE & IT Operations The Same Thing?](#)
- [The State of SRE Today](#)
- [ITOps vs DevOps vs NoOps: The IT Operations Evolution](#)

References

Adkins, Heather, Beyer, Betsy, Blankinship, Paul, Lewandowski, Piotr, Oprea, Ana, Stubblefield, Adam. *Building Secure & Reliable Systems*. O'Reilly Media, Inc., Sebastopol, CA, March 2020.

Beyer, Betsy, Jones, Chris, Murphy, Niall, Petoff, Jennifer. *Site Reliability Engineering*, <https://sre.google/sre-book/table-of-contents/>, O'Reilly, Accessed November 2021.

Beyer, Betsy, Kawahara, Kent, Murphy, Niall Richard, Rensin, David, Thorne, Stephen. *The Site Reliability Workbook*, <https://sre.google/workbook/table-of-contents/>, O'Reilly, Accessed November 2021.

Ewaschuk, Rob, SRE.Google Workbook: Chapter 6 – Monitoring Distributed Systems, <https://sre.google/sre-book/monitoring-distributed-systems/>, O'Reilly, Accessed November 2021.

Frame, Jess, SRE.Google Workbook: Chapter 4 – Monitoring, <https://sre.google/workbook/monitoring/>, Accessed November 2021.

Sloss, Benjamin Treynor, <https://www.oreilly.com/content/tenets-of-sre/>, October 6, 2017, Accessed November 2021.