



Avoiding the death of SRE documents that matter

April 2023

Yury Niño Roa

@yurynino
www.yurynino.dev

Cloud Infrastructure Engineer

 Google Cloud



October 4, 2018
Volume 16, Issue 4

PDF

Why SRE Documents Matter


How documentation enables SRE teams to manage new and existing services

Shylaja Nukala and Vivek Rau

SRE (site reliability engineering) is a job function, a mindset, and a set of engineering approaches for making web products and services run reliably. SREs operate at the intersection of software development and systems engineering to solve operational problems and engineer solutions to design, build, and run large-scale distributed systems scalably, reliably, and efficiently.

SRE core functions include:

- *Monitoring and metrics* — establishing desired service behavior, measuring how the service is actually behaving, and correcting discrepancies.
- *Emergency response* — noticing and responding effectively to service failures in order to preserve the service's conformance to its SLA (service-level agreement).
- *Capacity planning* — projecting future demand and ensuring that a service has enough computing resources in appropriate locations to satisfy that demand.
- *Service turn-up and turn-down* — deploying and removing computing resources for a service in a data center in a predictable fashion, often as a consequence of capacity planning.



Seeking
SRE

CONVERSATIONS ABOUT RUNNING PRODUCTION SYSTEMS AT SCALE

Agenda

Why SRE documents matter?


SRE Documents

How to keep live documents?

What Google learned?



A Context of SRE



Site Reliability Engineers operate at the intersection of software development and infrastructure engineering to solve operational problems and engineer solutions.




SRE Core Functions

- Monitoring and Metric.
- Emergency Response.
- Capacity Planning.
- Service turn-up and turn-down.
- Change Management.
- Performance.

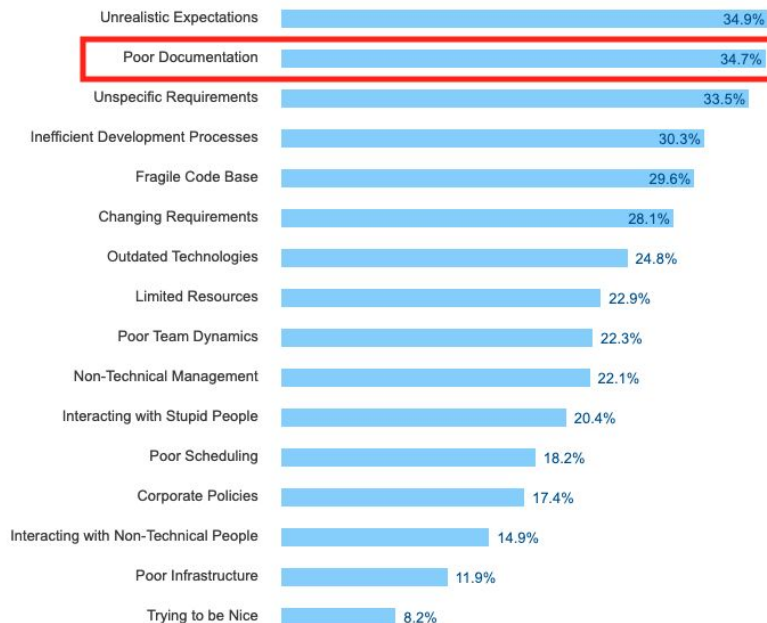
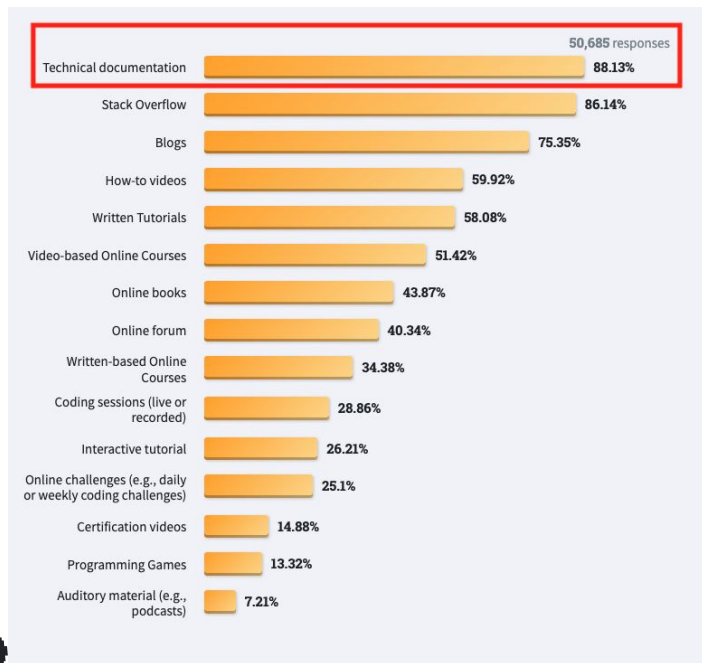
These requires bodies of documentation associated!



Regarding Documentation



VI. Challenges At Work



<https://survey.stackoverflow.co/2022/#overview>

Why SRE documents matter?



Because ...

If the **tribal knowledge** is not codified and documented, the concepts and principles will often need to be relearned **painfully** through trial and error.

Creating high-quality documentation that lays the foundation is a form that is easily discoverable, searchable, and maintainable.

New team members are trained through a systematic and well-planned induction and education program.




That is challenging ...

Time spent on documentation needs to come out of the development budget, and this is challenging

SREs often spend 35% of their time on operational work, which leaves only 65% for development.

Documentation is recognized or rewarded during performance review and promotion processes.



SRE Documents





SRE Documents ...



For New Service Onboarding

For Running a Service

For Production Products

For Reporting Service State

For Service Decommissioning


For Running SRE Teams



Documents for New Service Onboarding

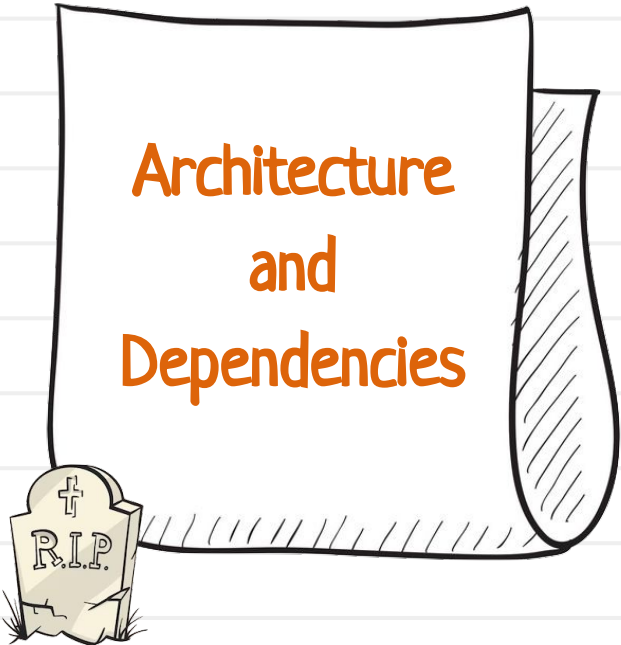
Production Readiness Review

A **PRR** (production readiness review) is conducted to make sure that a **service** meets accepted standards of operational readiness, and their **owners** have a **SRE guidance** about running them.



Docs for New Service Onboarding

Production Readiness Review



Architecture
and
Dependencies


- * What is your request flow from user to front end to back end?
- * Are there different types of requests with different latency requirements?

Docs for New Service Onboarding

Production Readiness Review

* How much traffic and rate of growth do you expect during and after the launch?

* Have you obtained all the compute resources needed to support your traffic?



Docs for New Service Onboarding

Production Readiness Review



Failure
Modes

* Do you have any single points of failure in your design?

* How do you mitigate unavailability of your dependencies?

Docs for New Service Onboarding

Production Readiness Review

* Are any manual processes required to keep the service running?

* How are we automating these processes?



Processes and
Automation

Docs for New Service Onboarding

Production Readiness Review



External Dependencies

- * What third-party code, data, services, or events do the service or the launch depend upon?
- * Do any partners depend on your service? If so, do they need to be notified of your launch?

Docs for New Service Onboarding

Explain the SRE role and responsibilities to set stakeholders expectations correctly.

Ensure that developer teams do not equate SREs with an Ops team.



SRE Role and
Responsibilities

Docs for New Service Onboarding




Engagement
Model
Document


- Service takeover criteria.
- SLO & Error budgets.
- New launch and launch freeze criteria.
- Service status reports.
- SRE staffing requirements.
- Feature roadmap planning process.

Documents for Running a Service

Running Service Documents are core operational assets SRE teams rely on to perform production services include service overviews, playbooks and procedures, postmortems, policies, and SLAs.




Docs for Running a Service



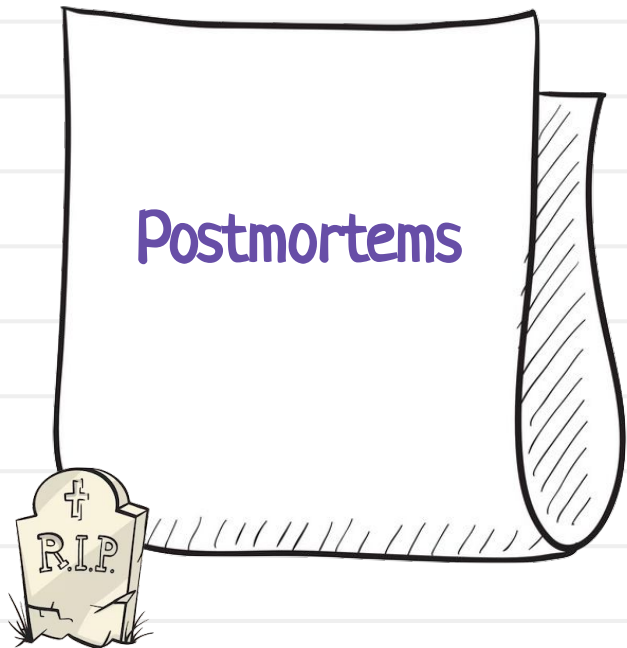
- * SREs need documents with enough information about a service to dig deeper.
- * This document provide a thorough description of the service and how it interacts with the world around it.

Docs for Running a Service

- * With the **playbooks**, the oncall engineers respond the alerts generated by service monitoring.
- * They contain commands and steps to review for accuracy.



Docs for Running a Service




Postmortems are an analysis conducted after a system failure:


- Timeline.
- Description of user impact.
- Root cause.
- Action items / lessons learned.

Docs for Running a Service

- * Technical Policies
- * Process Policies
- * Escalation Policies
- * Oncall Policies



Docs for Running a Service




* An **SLA** is a formal agreement with a customer on the performance a service commits to provide and what actions will be taken if that obligation is not met.


Documents for Production Products

Production Products Documents enable users to find out whether a product is right for them to adopt, how to get started, and how to get support.

They also provide a consistent user experience and facilitate product adoption.



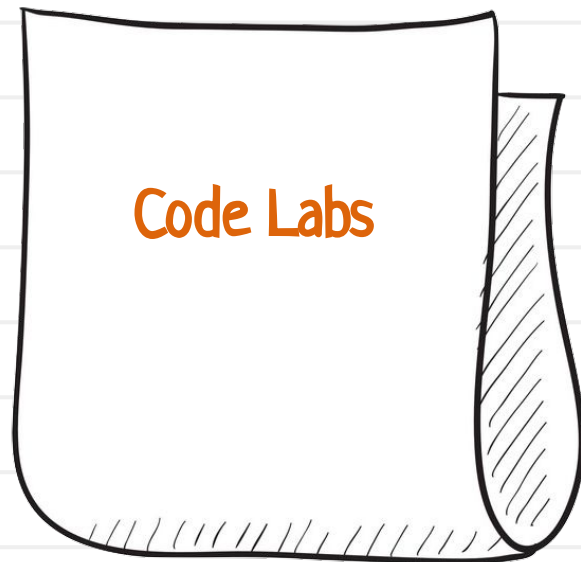
Docs for Production Products




- * Concepts Guide
- * Quickstart Guide
- * How-to Guide
- * Quickstart Guide
- * Developer Guide

Docs for Production Products

- * **Codelabs** provide in-depth scenarios that walk engineers step by step through a series of key tasks.
- * Engineers combine explanation, example code, and code exercises to get up to speed with the product.



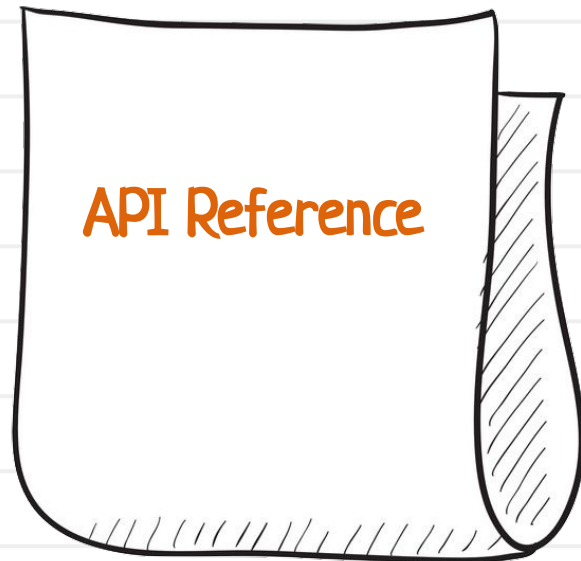
Docs for Production Products



- * The **FAQ page** answers common questions and covers caveats that users should be aware of.
- * **Support** page identifies how engineers can get help when they are stuck on something.

Docs for Production Products


API Reference provides descriptions of functions, classes, and methods, typically with minimal narrative or reader guidance.



Documents for Reporting Service State

This part describes the documents that SRE teams produce to communicate the state of the services they support.

That basically are: quarterly service review and a presentation about this.



Docs for Reporting Service State



Quarterly
Service Review
+
Presentation


A hand-drawn scroll with orange text. At the bottom left of the scroll is a small, hand-drawn tombstone with a cross on top and the letters 'R.I.P.' on it. The scroll has a wavy, ribbon-like edge on the right side.

The goal of a **quarterly report** is to cover a state of the service review, including details about performance, sustainability, risks, and overall production health.


Documents for Running SRE Teams

SRE teams need to have a cohesive set of reliable, discoverable documentation to function effectively as a team.

Some documents include: a Team Site and a Team charter



Docs for Running SRE Teams




- * A Team Charter explains the rationale for the team and documents its current major engagements.
- * A charter serves to establish the team identity, primary goals, and role relative to the rest of the organization.

Documents for New SRE Onboarding

SRE teams invest in training materials and processes for new SREs because training results in faster onboarding to the production environment.

Many SRE teams use checklists for oncall training.



Docs for Running SRE Teams



- * An **Oncall Checklist** covers all the high-level areas team members should understand well.
- * Examples include production concepts, front-end and back-end stack, automation and tools, and monitoring and logs.

Docs for Running SRE Teams



Role-Play
Trainings

* A classical example of this is the **Wheel of Misfortune** exercise, which presents an outage scenario to the team, with a set of data and signals that the hypothetical oncall SRE will need to use as input to resolve the outage.

How to keep live documents



Communicate the Value of Documentation

If you want to convince about documentation, it's essential that you demonstrate the quality, effectiveness, and value of your assets.


When you talk about the impact of your doc work, functional data is convincing.



Create a Repository

SRE team information can be scattered across a number of sites, local team knowledge, and Google Drive folders, which can make it difficult to find correct and relevant information.

A consistent structure will help team members find information quickly.



Create Templates

They make it easy for authors to create documentation by providing a clear structure that they can populate quickly with relevant information.

Templates make documentation easier to create and far easier to use.



Create Templates



Title

The title should be the name of the alert (e.g., Generic Alert_AlertTooGeneric).

Author:

Last updated:

Overview

Address the following:

- What does this alert mean?
- Is it a paging or an email-only alert?
- What factors contributed to the alert?
- What parts of the service are affected?
- What other alerts accompany this alert?
- Who should be notified?

Alert Severity

Indicate the reason for the severity (email or paging) of the alert and the impact of the alerted condition on the system or service.

Verification

Provide specific instructions on how to verify the condition.

Troubleshooting

List and describe techniques and related links to relevant dashboards. Include wa



Define Success Metrics

As you define your documentation requirements, it's also important to define **how you will measure the functional** quality of your docs.

For example a service overview has high impact if its usage is measured and the times of solving an incident are reduced them.



Follow Tech writing Practices

It is important to have guidance from technical writers on best practices for working with SRE teams. They should partner with SREs to provide operational documentation for running services and product documentation for SRE products and features.



Require Docs as Part of Code Review

Here's a good rule of thumb: Doing Docs Better: Best Practices!

If a developer, SRE, or user of your project needs to change their behavior after this change, the changelist should include doc changes.





Thank you so much!