

ACADEMICIA

ISSN (online) : 2249-7137

## ACADEMICIA

An International  
Multidisciplinary Research  
Journal



Published by  
**South Asian Academic Research Journals**  
A Publication of CDL College of Education, Jagadhri  
(Affiliated to Kurukshetra University, Kurukshetra, India)

**ACADEMICIA**

An International Multidisciplinary Research Journal

ISSN (online) : 2249 –7137

Editor-in-Chief : Dr. B.S. Rai

Impact Factor : SJIF 2020 = 7.13

Frequency : Monthly

Country : India

Language : English

Start Year : 2011

Indexed/ Abstracted : Scientific Journal Impact Factor (SJIF2020 - 7.13), Google Scholar, CNKI Scholar, EBSCO Discovery, Summon (ProQuest), Primo and Primo Central, I2OR, ESJI, IJIF, DRJI, Indian Science and ISRA-JIF and Global Impact Factor 2019 - 0.682


E-mail id: saarjournal@gmail.com

**VISION**

The vision of the journals is to provide an academic platform to scholars all over the world to publish their novel, original, empirical and high quality research work. It propose to encourage research relating to latest trends and practices in international business, finance, banking, service marketing, human resource management, corporate governance, social responsibility and emerging paradigms in allied areas of management including social sciences , education and information & technology. It intends to reach the researcher's with plethora of knowledge to generate a pool of research content and propose problem solving models to address the current and emerging issues at the national and international level. Further, it aims to share and disseminate the empirical research findings with academia, industry, policy makers, and consultants with an approach to incorporate the research recommendations for the benefit of one and all.



**ACADEMICIA**  
An International  
Multidisciplinary  
Research Journal  
(Double Blind Refereed & Reviewed International Journal)



**SOUTH ASIAN ACADEMIC RESEARCH  
JOURNALS ([www.saarj.com](http://www.saarj.com))**

**ACADEMICIA: An International Multidisciplinary  
Research Journal**

**ISSN: 2249-7137 Impact Factor: SJIF 2022 = 8.252**

**SPECIAL ISSUE ON  
"ENGINEERING APECTS OF SITE  
RELIABILITY"**

**July 2022**



# ACADEMICIA

## An International Multidisciplinary Multidisciplinary Research Journal

(Double Blind Refereed & Reviewed International Journal)



SR. NO.	PARTICULAR	PAGE NO
1.	<b>INTRODUCTION TO SITE RELIABILITY ENGINEERING</b> Mr. Medikeranahalli Santhosh	<b>6-14</b>
2.	<b>INTERSECTION OF SECURITY AND RELIABILITY</b> Ms. Anju Mathew	<b>15-23</b>
3.	<b>A BRIEF INTRODUCTION ON UNDERSTANDING ADVERSARIES</b> Ms. Appaji Gowda Shwetha	<b>24-32</b>
4.	<b>A CASE STUDY ON SAFE PROXIES</b> Mr. Bhavan Kumar	<b>33-38</b>
5.	<b>A BRIEF INTRODUCTION ON DESIGN TRADEOFFS</b> Dr. Nakul Ramanna Sanjeevaiah	<b>39-46</b>
6.	<b>DESIGN FOR LEAST PRIVILEGE</b> Dr. Shrishail Anadinni	<b>47-53</b>
7.	<b>AUTHORIZATION DECISIONS, TRADEOFFS AND TENSIONS</b> Ms. Hireballa Sangeetha	<b>54-61</b>
8.	<b>DESIGN FOR UNDERSTANDABILITY</b> Mr. Jayaraj Dayalan	<b>62-70</b>
9.	<b>SYSTEM ARCHITECTURE AND SOFTWARE DESIGN</b> Dr. Ganpathi Chandankeri	<b>71-79</b>
10.	<b>DESIGN FOR A CHANGING LANDSCAPE</b> Mr. Narayana Gopalakrishnan	<b>80-89</b>
11.	<b>DESIGN FOR RESILIENCE</b> Dr. Jagdish Godihal	<b>90-97</b>
12.	<b>CONTROLLING DEGRADATION AND BLAST RADIUS</b> Dr. Mohammad Shahid Gulgundi	<b>98-107</b>

13.	<b>FAILURE DOMAINS AND REDUNDANCIES</b> Mr. Ahamed Sharif	108-116
14.	<b>DESIGN FOR RECOVERY</b> Ms. Aashi Agarwal	117-125
15.	<b>MITIGATING DENIAL-OF-SERVICE ATTACKS</b> Dr. Ganpathi Chandankeri	126-133
16.	<b>CASE STUDY: DESIGNING, IMPLEMENTING, AND MAINTAINING A PUBLICLY TRUSTED CA</b> Mr. Narayana Gopalakrishnan	134-141
17.	<b>A STUDY TO EVALUATING AND BUILDING FRAMEWORKS</b> Dr. Jagdish Godihal	142-149
18.	<b>A BRIEF INTRODUCTION ON SECURITY AND RELIABILITY</b> Dr. Mohammad Shahid Gulgundi	150-155
19.	<b>A STUDY ON TYPES OF TESTING CODES</b> Mr. Ahamed Sharif	156-163
20.	<b>CONCEPT OF DEPLOYING CODES</b> Ms. Aashi Agarwal	164-172
21.	<b>INVESTIGATING SYSTEMS AND SECURE DEBUGGING ACCESS</b> Dr. Topraladoddi Madhavi	173-179
22.	<b>BUILDING A CULTURE OF SECURITY AND RELIABILITY</b> Dr. Shrishail Anadinni	180-186

## INTRODUCTION TO SITE RELIABILITY ENGINEERING

**Mr. Medikeranahalli Santhosh\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:santhoshmb@presidencyuniversity.in

---

### ABSTRACT:

*The main ideas and tenets of SRE are briefly discussed in this abstract, with an emphasis on how it promotes operational excellence and reduces system failures. The main duties of an SRE team are examined, including as monitoring, incident response, capacity planning, and performance optimization. The vital role that automation plays in SRE, which makes it possible to manage infrastructure effectively while reducing human error, is also covered in the paper. The paper also explores the organizational and cultural components of SRE, highlighting the value of cooperation, communication, and ongoing development. It emphasizes the use of reliability engineering techniques across the full lifespan of software development, from planning and creation through deployment and continuous maintenance. The paper ends by recognising SRE's influence on contemporary technology organisations, including how it encourages a culture of dependability, facilitates quick innovation, and raises customer satisfaction. SRE equips teams to create and manage robust systems that can resist the demands of a dynamic and ever-changing technological world by fusing software engineering and operations skills.*

**KEYWORDS:** Incident Management, Replication, Robust System, Site Reliability.

---

### INTRODUCTION

High availability, performance, and resilience are essential in the contemporary digital environment where organisations significantly depend on online platforms and services. SRE tackles these issues by encouraging a proactive, methodical method of managing complex systems and by integrating engineering ideas into operations. SRE is fundamentally motivated by the aim to balance innovation and stability. It highlights the importance of service dependability while facilitating quick feature development and deployment. This strategy encourages cooperation between the development and operations teams in order to accomplish common objectives. It is consistent with the wider DevOps concept[1]–[3].

Site Reliability Engineering's guiding concepts include the following:

1. SLOs (Service Level Objectives)
2. Automation
3. The Monitoring and Alerting Process
4. Incident Management
5. Capacity Planning

## 6. Cross-Functional Collaboration

### DISCUSSION

Software engineering and operations are combined in the field of site reliability engineering (SRE), which helps create and manage highly dependable systems. To manage their extensive, complex infrastructure and services, Google created it. SRE prioritizes the creation of reliable, scalable, and effective systems while preserving a balance between feature development and dependability. The main components and tenets of site reliability engineering are as follows:

#### Reliability

Site Reliability Engineering (SRE) is based on the basic concept of reliability. Reliability in SRE refers to a system's or service's capacity to continuously carry out its intended function under anticipated circumstances while satisfying specified service level agreements (SLAs).

The following are some essential elements of site reliability engineering:

- a. **Service Level Objectives (SLOs):** SLOs, which are precise goals for a service's dependability, are defined and established by SREs in collaboration with stakeholders. Availability, latency, error rates, and other pertinent performance characteristics are often included in SLOs. These goals help create expectations and give dependability a quantifiable goal.
- b. **Error Budgets:** To strike a balance between dependability and feature development, SREs leverage the idea of error budgets. The allowed amount of downtime or service degradation during a certain time period is represented by an error budget. SREs and development teams may decide the trade-offs between stability and adding new features or changing the system by establishing an error budget.
- c. **Alerting and Monitoring:** Systems for monitoring are essential for assuring dependability. SREs develop thorough monitoring frameworks to gather and examine important metrics and system health indicators. This enables them to see abnormalities, identify future problems, and take proactive action to preserve or restore dependability. SREs are informed via alerting systems when predetermined thresholds are crossed, allowing them to respond appropriately right away.
- d. **Incident Management:** In order to lessen the effects of service failures or interruptions, SREs place a major emphasis on incident management. They set up systems for detecting, escalating, and resolving incidents as part of incident response. When accidents happen, SREs collaborate with development teams to quickly find solutions, record the specifics, and carry out post-incident evaluations to draw lessons from the experience and avert future occurrences.
- e. **Fault Tolerance and Redundancy:** To achieve high availability, SREs design systems with fault tolerance and redundancy in mind. Implementing techniques like replication, load balancing, failover, and backup plans is required for this. SREs improve system dependability by spreading the workload over many components or data centers and removing single points of failure.

- f. **Proactive Maintenance and Capacity Planning:** SREs plan their capacity and do proactive maintenance tasks to foresee and resolve possible reliability problems. They do capacity evaluations, historical data analysis, routine system audits, and demand forecasting. This aids SREs in identifying and reducing risks related to scalability, performance, and resource limitations before they have an effect on dependability.
- g. **Testing and Release Practices:** SREs promote rigorous testing and release procedures in order to preserve dependability. They design efficient testing methodologies, such as unit tests, integration tests, and performance tests, in close collaboration with development teams. In order to make sure that new features and updates are thoroughly tested and validated before being introduced to production environments, SREs also work with developers to establish and implement safe and dependable deployment methods.
- h. **Continuous Improvement:** In SRE, reliability is a constant endeavor. SREs gather input, regularly assess system performance, and pinpoint opportunities for development. To increase system dependability, they refine procedures, boost settings, and put new features into practice. This cycle of continuous improvement must include information sharing across teams, performing blameless postmortems, and learning from occurrences.

SRE assists organisations in developing and maintaining systems that provide reliable and consistent services by placing a strong emphasis on dependability. SREs work to achieve and surpass the established reliability objectives while assuring the overall performance of the systems they support via proactive measures, monitoring, incident management, fault tolerance, and continuous improvement.

### Automation

A key component of Site Reliability Engineering (SRE) is automation. It is essential for enhancing the scalability, efficiency, and dependability of the system. The following are some significant SRE applications for automation: Management of Software Deployment and Configuration: SREs automate software system deployment and configuration. To design and manage the intended state of infrastructure and applications, they employ tools like configuration management systems (e.g., Puppet, Chef, Ansible) or infrastructure-as-code frameworks (e.g., Terraform). Automation promotes consistency, minimizes human error, and permits dependable and quick deployments[4]–[6].

SREs may decrease human work, boost productivity, and boost dependability by automating these numerous facets of system administration and operations. Automation makes processes quicker and more reliable, lowers the possibility of human mistake, and frees up SREs to concentrate on more strategic and beneficial tasks.

### Incident Management

A key component of Site Reliability Engineering (SRE) is incident management. Effectively managing and mitigating events is the responsibility of SREs in order to reduce their negative effects on system availability and performance. The essential elements of incident management in SRE are as follows:

**Incident Response:** To respond to events quickly and effectively, SREs adhere to defined incident response procedures. When an event happens, they move swiftly to determine its origin,



lessen its effects, and resume regular service. An incident response team must be put together, communication lines must be coordinated, and predetermined response processes must be followed.

- a. **Incident Escalation:** SREs have a clear escalation mechanism to make sure the right people or teams are handling events. They classify issues according to their effect, severity, and urgency, and when required, they escalate them to higher-level support or engineering teams. Effective escalation protocols aid in quick resolution and stop events from turning into significant disruptions.
- b. **Incident Communication:** During incident management, communication must be precise and timely. SREs keep lines of communication open with all relevant parties to provide updates on the incident's progress and anticipated time of resolution, including customers, support teams, and management. To communicate information and control stakeholder expectations, they employ incident communication channels including status pages, incident response tools, and alerts.
- c. **Post-Incident Analysis:** SREs undertake post-event analysis, often known as postmortems, after an incident has been resolved in order to identify the underlying causes, contributing factors, and lessons learned. This entails a blameless analysis that emphasises system improvements above personal accountability. SREs record the specifics of the occurrence, communicate the results to the rest of the team, and promote improvements to stop similar accidents from happening again.
- d. **Continuous Improvement:** SRE incident management places a strong emphasis on a culture of ongoing development. SREs aggressively seek for methods to improve system resilience and dependability by learning from accidents. They identify reoccurring problems, suggest system or process upgrades, and work with development teams to put preventative measures into practice. Continuous improvement lessens the probability and severity of future events while strengthening the system's overall dependability.
- e. **Incident Documentation:** SREs keep thorough records of incidents, such as run books, postmortems, and incident reports. These records outline the occurrence, the steps followed, and the resolution procedure. In addition to aiding in the training of new team members and supporting incident response activities by giving historical context, event recording acts as a knowledge foundation for future use.
- f. **Incident Management Tools:** To make their incident response operations more efficient, SREs use a variety of incident management tools and technology. These resources include of communication channels, collaboration platforms, monitoring and alerting systems, and incident tracking systems. Within the incident response team, good coordination, real-time event tracking, and effective communication are made possible by incident management technologies.

SREs guarantee that problems are immediately handled, that their effects are reduced, and that system dependability is continuously increased by adhering to certain incident management practices. Maintaining high availability, minimizing downtime, and delivering a great user experience all depend on effective incident management.

### Monitoring and Alerting Process

---

A strong Site Reliability Engineering (SRE) practise must include monitoring and alerting. They make it possible for proactive problem detection, quick incident response, and timely resolution. An outline of the monitoring and alerting procedure is provided below:

- a. **Define the monitoring goals:** Establish definite monitoring goals that are in line with the important KPIs, performance measures, and user experience of the system. The intended degree of system availability, performance, and dependability should be reflected in these goals. Think about elements such as reaction time, error rates, throughput, and resource use.
- b. **Identify Key Metrics:** Choose the important metrics that provide information on the behaviour and condition of the system. In addition to request counts and error rates, these metrics may also include CPU and memory consumption, network traffic, database slowness, and others. Choose metrics that have a direct influence on the functionality, dependability, and user experience of the system.
- c. **Implement Monitoring Tools:** Select monitoring programs and solutions that enable the gathering, storing, visualizing, and analysing of metrics. Prometheus, Grafana, Datadog, New Relic, and Nagios are a few examples of popular tools. To begin gathering pertinent data, integrate these technologies with your infrastructure and apps[7]–[9].
- d. **Set Alerting Thresholds:** Set thresholds or requirements for notifications depending on the chosen metrics. These thresholds are predetermined numbers that show when a measure behaves differently from what is expected or within acceptable bounds. Alerting thresholds may be computed dynamically using historical data or statistical techniques, or they might be static numbers.
- e. **Configure Alerting Mechanisms:** Set up alerting systems to inform the necessary parties when thresholds are crossed or abnormalities are found. Numerous techniques are used often, such as email notifications, SMS alerts, platforms for instant messaging, and interaction with event management tools like PagerDuty or OpsGenie. Make that the appropriate teams or individuals who are in charge of incident response get notifications.
- f. **Incident Response:** Follow established incident response protocols when an alert is triggered. Assign team members participating in incident resolution positions and duties. Investigate the problem right away, determine its source, and take the appropriate steps to lessen its effects. Keep a record of the incident's specifics for future learning and reference.

### Capacity Planning

The process of figuring out the resources and infrastructure needed to fulfil present and future workload needs is known as capacity planning. To maintain optimum system performance and availability, it entails analysing previous data, projecting future development, and making knowledgeable judgements regarding the allocation of resources. Site Reliability Engineering (SRE) must include capacity planning in order to avoid service interruptions brought on by resource shortages.

The following are the key steps in capacity planning:

- a. **Requirements for gathering:** Gathering needs from multiple stakeholders, such as company owners, product managers, and development teams, is the first stage in capacity planning.

Understanding the anticipated workload, user traffic patterns, expected response times, and any particular performance or scalability targets are all part of this.

- b. **Baseline Analysis:** In order to create a baseline for system behavior, SREs examine historical data. This entails monitoring metrics over time, including CPU utilization, memory use, network traffic, and storage capacity. The identification of patterns, peak times, and trends in resource usage is aided by baseline analysis.
- c. **Forecasting:** SREs assess future resource needs based on past information and projected changes in workload. This may include doing statistical analysis, predicting trends, or using particular forecasting models. The objective is to foresee the increase in demand and allocate resources appropriately.
- d. **Sizing and Resource Allocation:** SREs choose the right resource size and distribution to meet the anticipated workload. This involves assessing the needs for the network, CPU, memory, and storage. Peak traffic, projected growth rates, and performance goals are just a few examples of the variables that might influence sizing. SREs work with infrastructure teams to efficiently distribute resources[10].
- e. **Performance Validation and Testing:** To make sure that the resources assigned are enough to manage the projected demand, SREs do performance testing and validation. This include testing the system under different loads, replicating real-world circumstances, and analysing system performance. Performance testing findings assist in validating the capacity plan and identifying any bottlenecks or restrictions.
- f. **Alerting and Monitoring:** After the capacity plan is put into place, SREs create monitoring and alerting systems to keep track of important performance indicators and make sure the system is operating at the projected capacity. Monitoring identifies resource consumption variances and issues alarms when predetermined thresholds are crossed.
- g. **Elasticity and Scalability:** SREs consider elasticity and scalability while designing systems. This makes it possible for the infrastructure to dynamically modify its capacity in response to current demand. Vertical scaling involves adding resources to the current infrastructure; horizontal scaling involves adding new instances or nodes. Elasticity makes ensuring the system can automatically scale up or down in response to changes in workload.
- h. **Capacity Management:** The system's capacity is continually managed and optimized by SREs. This entails keeping track of how resources are being used, examining performance information, and pinpointing areas that may be improved. To make the best use of resources, carry out performance improvements, and resolve any capacity issues, SREs work with development teams.
- i. **Regular Review and Iteration:** Iterative processes are used in capacity planning. SREs evaluate the efficacy of the capacity plan on a regular basis and make necessary revisions. To make sure the system can accommodate changing needs, the capacity plan should be revised as the workload and business requirements change.

To guarantee that systems can manage the anticipated demand and provide the appropriate level of performance and availability, capacity planning is essential. SREs may avoid service

interruptions brought on by capacity restrictions and guarantee a positive user experience by proactively analysing resource needs and making knowledgeable choices.

### **Disaster Recovery**

Site Reliability Engineering (SRE), which focuses on limiting the effects of catastrophic events that might interrupt or harm systems and services, includes disaster recovery as a key component. Catastrophe recovery in SRE aims to reduce customer impact, downtime, and data loss in the case of a catastrophe. The following are the salient features of SRE's catastrophe recovery:

- a. **Risk Assessment:** To find probable catastrophes and their potential effects on systems and services, SREs undertake an extensive risk assessment. This includes calamities caused by nature, issues with hardware or software, cyberattacks, and other possible dangers. The evaluation helps in setting priorities and creating a catastrophe recovery strategy.
- b. **Business Impact Analysis:** To determine the importance and priority of various systems and services, SREs do a business impact study. For each system or service, this analysis aids in determining the recovery goals, recovery time objectives (RTO), and recovery point objectives(RPO). The maximum allowable downtime is defined by the RTO, while the maximum allowable data loss is defined by the RPO.
- c. **Planning for Disaster Recovery:** The stages, methods, and procedures for recovering systems and services in the case of a catastrophe are laid out in a thorough disaster recovery plan that SREs create. The strategy outlines the team members' roles and duties, communication guidelines, and particular recovery techniques for certain circumstances.
- d. **Fault-Tolerant System Design:** To lessen the effects of catastrophes, SREs include fault tolerance into the architecture of their systems. This entails putting redundancy into practise at several levels, including using numerous data centres, redundant hardware, and data replication. Redundancy makes ensuring there is a backup system or location that can take over immediately if one component or location fails.
- e. **Backup and Restore:** SREs often use backup techniques to protect data integrity and speed up recovery. This entails regularly backing up important data and settings and keeping backups in safe offsite places. Periodic restoration tests are also carried out by SREs to confirm the accuracy and accessibility of the backup data.
- f. **Replication and Failover:** Replication methods are used by SREs to keep current copies of the data and services in many locations or data centers. In the case of a catastrophe, systems may easily transfer to a backup location or instance thanks to replication's enabled failover capabilities. Data loss and downtime are reduced as a result.
- g. **Disaster Recovery Drills:** SREs regularly assess the efficacy of the recovery plan and look for any holes or problems through drills or exercises for disaster recovery. The team's responsiveness, recovery time, and data integrity are assessed throughout these exercises, which replicate numerous catastrophe situations. The disaster recovery strategy is improved and refined using the lessons learnt from these simulations.
- h. **Communication and Incident Management:** Effective communication is vital during a crisis. To guarantee quick and accurate communication among team members, stakeholders,

and clients, SREs build communication standards and channels. Any disaster-related occurrences are managed and their lessons learned using incident management techniques including incident response and post-incident analysis.

- i. **Continuous Improvement:** Based on the lessons acquired from training exercises, actual disasters, and changing business needs, SREs regularly review and enhance the disaster recovery plan. To increase the resiliency and efficiency of the disaster recovery process, they use industry best practices and keep current on new threats and technology.

The goal of SRE disaster recovery is to lessen the effects of catastrophes and guarantee the continuation of systems and services. SREs can quickly recover from catastrophes, minimise downtime, and safeguard the integrity and availability of crucial systems and data by anticipatorily designing, deploying redundancy, and routinely testing the recovery processes. In order to create and manage robust systems, site reliability engineering integrates software engineering, operations, and a reliability mindset. In order to offer dependable and scalable services while promoting cooperation between various teams engaged in the creation and maintenance of systems, it places a strong emphasis on preventative measures, automation, and continuous improvement.

## CONCLUSION

SRE (Site Reliability Engineering), a crucial field that focuses on assuring the dependability, availability, and performance of software systems and services, is summarised above. Designing, constructing, and maintaining reliable and scalable systems requires the use of software engineering concepts and operational knowledge. As failure is unavoidable and should be controlled pro-actively rather than reactively, the major objective of SRE is to find a balance between dependability and innovation. SRE teams seek to decrease downtime, issues, and bad user experiences by adopting a data-driven and automation-centered strategy. SRE is not a universally applicable solution, however. Organizational commitment, cultural changes, and continuous investment in automation and tooling are all necessary. SRE is a journey that calls for constant learning, adaptation, and iteration to stay up with changing business demands and technological requirements. In conclusion, Site Reliability Engineering is essential to contemporary software systems because it enables businesses to provide dependable, scalable, and high-performing services. Organisations may handle operational issues and gain a competitive edge in the modern digital environment by using SRE concepts.

## REFERENCES

- [1] P. Coussement, J. Maertens, J. Beauprez, W. Van Bellegem, and M. De Mey, "One step DNA assembly for combinatorial metabolic engineering," *Metab. Eng.*, 2014, doi: 10.1016/j.ymben.2014.02.012.
- [2] N. Poltorachenko, "Simulation Of The Initial Stage Of Engineering Network Design," *Manag. Dev. Complex Syst.*, 2021, doi: 10.32347/2412-9933.2021.45.97-101.
- [3] J. C. Langford, "Application of reliability methods to the design of underground structures," *ProQuest Diss. Theses*, 2013.
- [4] M. Rezaei, Y. Li, X. Li, and C. Li, "Improving the Accuracy of Protein-Ligand Binding Affinity Prediction by Deep Learning Models : Benchmark and Model," *ChemRxiv*, 2019.

- [5] J. Sung-Gu, "Wearable device in security," *Int. J. Secur. its Appl.*, 2015, doi: 10.14257/ijasia.2015.9.6.23.
- [6] Y. Bi *et al.*, "An asymmetric PCR-based, reliable and rapid single-tube native DNA engineering strategy," *BMC Biotechnol.*, 2012, doi: 10.1186/1472-6750-12-39.
- [7] M. Wathelet, "Array recordings of ambient vibrations : surface-wave inversion," *Liège Univ.*, 2005.
- [8] P. L. McCarty, "Engineering concepts for in situ bioremediation," *J. Hazard. Mater.*, 1991, doi: 10.1016/0304-3894(91)87002-J.
- [9] B. Valpy, G. Hundleby, K. Freeman, A. Roberts, and A. Logan, "Future renewable energy costs: offshore wind," 2017.
- [10] G. G. Kashevarova, Y. L. Tonkov, and I. I. Tonkov, "Intellectual Automation Of Engineering Survey Of Building Objects," *Int. J. Comput. Civ. Struct. Eng.*, 2017, doi: 10.22337/1524-5845-2017-13-3-42-57.

## INTERSECTION OF SECURITY AND RELIABILITY

**Ms. Anju Mathew\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id: anjumathew@presidencyuniversity.in

---

### ABSTRACT:

*Privacy and security are two notions that go hand in hand. A system must be fundamentally secure and operate as intended in the face of an attacker in order to protect user privacy. Similar to this, if a system doesn't respect user privacy, it won't satisfy the demands of many users. Although the emphasis of this book is security, achieving privacy goals may often be accomplished using the broad strategies we provide. The means of cooperation and the information that is accessible to responders during an event may both be impacted by the existence of an adversary. Responders to reliability events who can swiftly identify and address the main cause and have a variety of viewpoints are advantageous. To prevent the adversary from learning about the recovery process, you should typically manage security issues with the fewest amount of individuals possible. You will only provide information to those who require it in the security situation. Similar to this, extensive system logs may speed up recovery time and help with incident response, but depending on what is documented, those logs might also be a lucrative target for an attacker.*

**KEYWORDS:** Availability, Confidentiality, Investigating Systems, Resilience.

---

### INTRODUCTION

The ideas of security and dependability have merged in today's digital world, where technology plays a crucial part in our daily lives. Protecting systems and data against cyber-attacks while maintaining the continuous availability and performance of their services is a crucial problem faced by organisations across sectors. This fusion of security and dependability lays the groundwork for a comprehensive and effective strategy for addressing technology risks. Protecting systems, networks, and data against unauthorized access, breaches, and harmful actions is referred to as security. It includes several different fields of study, including network security, data encryption, access restrictions, vulnerability management, and incident response. The capacity of systems and services to consistently provide their intended functionality in favorable and unfavorable circumstances is the emphasis of dependability, on the other hand. High availability must be guaranteed, downtime must be kept to a minimum, and faults or disturbances must be dealt with quickly [1]–[3].

Organisations work to create a harmonic balance at the vital intersection of security and dependability. Prioritizing one component above another is no longer adequate. While a dependable system with insufficient security measures may be vulnerable to catastrophic breaches and compromises, a highly secure but unreliable system may still be prone to substantial disruptions.

Organisations may successfully reduce risks and improve their overall resilience by combining security and dependability into a cohesive strategy. Implementing strong security measures that are smoothly woven into the structure of dependable systems is part of this integrated strategy. It entails keeping an eye on security as a constant concern during the design, development, and deployment phases as well as continually modifying security protocols to handle new threats. Beyond technical considerations, security and dependability are intertwined. Organisations must promote a security-conscious culture where staff members are aware of the value of security procedures and actively involved in safeguarding systems and data. Additionally, legal and regulatory compliance are crucial for maintaining security and dependability since organisations are required to follow them in order to protect sensitive data and keep stakeholders' confidence.

The nexus of security and dependability is more important than ever in this age of more sophisticated cyber-attacks. To address these issues and create resilient systems that can resist the constantly changing threat environment, organisations must take a proactive, all-encompassing, and holistic strategy. Organisations may inspire trust in their clients, partners, and stakeholders and preserve a competitive advantage in the digital world by efficiently managing risks, maintaining the availability and integrity of systems, and protecting sensitive data.

## **DISCUSSION**

A really trustworthy system must have both reliability and security, but it is challenging to create systems that have both. Although the criteria for security and dependability have many similar characteristics, they also call for various design considerations. It is simple to overlook the delicate interactions between security and dependability that might result in unanticipated results. Poor load-balancing and load-shedding tactics caused a dependability issue that led to the password manager's collapse, and several security-related precautions made it more difficult for it to recover.

### **Reliability versus Security: Design Considerations**

You must take into account several hazards while planning for dependability and security. The main dependability concerns are not intentional, such a faulty software update or a broken physical item. However, security concerns originate from adversaries who are actively attempting to exploit system flaws. When you design for dependability, you make the assumption that certain things will eventually go wrong. When creating security-conscious designs, you must consider the possibility that an enemy may attempt to cause problems at any time.

As a consequence, many systems are built to react to failures in quite distinctive ways. Systems often fail safe (or open) in the absence of an adversary; for instance, an electronic lock is built to stay open in the event of a power outage, allowing safe escape via the door. Security flaws that are evident may result from fail safe or open behavior. You may design the door to fail secure and stay closed when the power is out to protect yourself from an enemy who could take advantage of a power outage.


### **Confidentiality, Integrity, Availability**



The confidentiality, integrity, and availability of systems are issues that both security and dependability are interested in, although they approach these qualities from different angles. The existence or absence of a hostile enemy is the primary distinction between the two points of view. A trustworthy system must not unintentionally violate confidentiality, unlike a chat programme with bugs that misdelivers, garbles, or loses messages. A secure system must also guard against active adversaries accessing, altering, or erasing sensitive data. Let's look at a few instances that show how a dependability issue might result in a security vulnerability.

- 1. Confidentiality:** A prominent confidentiality issue in the aviation sector is a push-to-talk microphone that is locked in the broadcast position. A jammed microphone has exposed private communications between pilots in the cockpit in a number of well-documented instances, which is a violation of confidentiality. In this instance, there is no deliberate opponent at play since the gadget transmits when the pilot does not want it to due to a hardware dependability defect.
- 2. Integrity:** Similar to that, a breach of data integrity need not entail an active opponent. In 2015, Google's Site Reliability Engineers (SREs) discovered that certain blocks of data's end-to-end cryptographic integrity checks were failing. The SREs made the decision to create software that exhaustively calculated the integrity check for every version of the data with a single-bit flip (a 0 converted to a 1, or vice versa), since some of the computers that processed the data afterwards showed indications of unfixable memory problems. They may then check to see whether one of the outcomes matched the integrity check's initial value. In fact, all mistakes were single-bit flips, and the SREs were able to recover every bit of data. It's interesting that in this case, a security measure saved the day during a dependability crisis. (Google's storage systems also include noncryptographic end-to-end integrity checks, but additional problems made it impossible for SREs to identify the bit flips.)
- 3. Availability:** Finally, it goes without saying that availability raises reliability and security issues. An attacker could take advantage of a system's flaw to disable it or make it difficult for authorized users to utilise it. A distributed denial-of-service (DDoS) assault occurs when a large number of devices are controlled and deployed around the globe to overwhelm a target with traffic[4]–[6].

Attacks that induce a denial of service (DoS) provide a challenging situation since they compromise both security and dependability. A malicious attack could seem to a victim to be no different from a real surge in traffic or a design mistake. An unanticipated burden on Google's central time server resulted from a 2018 software upgrade that prompted certain Google Home and Chromecast devices to create significant synchronized network traffic surges when they reset their clocks. Similar to a classic application-level DDoS assault, a significant breaking news item or other event that drives millions of users to submit almost identical requests might also seem extremely similar. Figure 1 depicts the surge in searches that occurred after a 4.5 magnitude earthquake struck the San Francisco Bay Area in the middle of the night in October 2019.



**Figure 1: Web traffic, measured in HTTP requests per second, reaching Google infrastructure serving users in the San Francisco Bay Area when a magnitude 4.5 earthquake hit the region on October 14, 2019.**

### **Reliability and Security: Commonalities**

Contrary to many other system features, dependability and security are emergent qualities of a system's architecture. Both are also difficult to add after the fact, therefore you should preferably take both into consideration from the beginning of design. Because it is simple for system modifications to unintentionally damage them, they also need constant monitoring and testing over the whole system lifespan. A seemingly harmless update to one component may end up having an impact on the dependability or security of the whole system in a manner that may not be obvious until it causes an incident. In a complex system, reliability and security qualities are often defined by the interaction of several components. Let's look more closely at these and other similarities.

### **Invisibility**

When everything is running well, reliability and security are mostly unseen. But gaining and maintaining the confidence of clients and partners is one of the objectives of dependability and security teams. A strong basis for establishing trust is good communication, both in difficult times and when things are going well. The information must be as accurate, concrete, and clear of jargon and cliches as is humanly feasible. Unfortunately, dependability and security are often seen as expenses that may be cut or postponed without immediate repercussions due to their inherent obscurity in the absence of emergency. Failures in security and dependability may, nevertheless, be quite expensive. The amount Verizon paid to purchase Yahoo!' may have been reduced by \$350 million as a result of data breaches, according to media sources.'s online store in 2017. Almost 700 flights were cancelled and hundreds of others were delayed due to a power outage at Delta Airlines in the same year, which reduced the airline's daily flight throughput by almost 60%.

### **Assessment**

You may use risk-based methodologies to assess the costs of adverse occurrences, as well as the upfront and opportunity costs of averting these events, since it is not practicable to attain complete dependability or security. For dependability and security, you should estimate the likelihood of unfavorable occurrences differently. You may assume independence of failures across the various components, which allows you to reason about the reliability of a composition of systems and plan engineering work according to specified error budgets<sup>1</sup>. Such a composition's security is more challenging to evaluate. The design and execution of a system may be examined for some degree of confidence. Adversarial testing, which simulates assaults from the viewpoint of a predetermined opponent, may be used to assess a system's resilience to certain attack types, the efficacy of attack detection methods, and the possible repercussions of attacks.

### **Simplicity**

One of the greatest methods to enhance your capacity to evaluate both the security and dependability of a system is to keep system design as simple as feasible. A more straightforward design minimizes the attack surface, lowers the possibility of unexpected system interactions, and makes the system simpler for people to understand and reason about. Understanding is particularly helpful in times of crisis since it enables responders to lessen symptoms rapidly and shorten the mean time to repair (MTTR). In further detail, Chapter 6 goes through methods for reducing attack surfaces and dividing up the responsibility for security invariants into manageable, basic subsystems that can be thought about separately.

### **Evolution**

Despite their initial simplicity and elegance, systems seldom endure throughout time unmodified. Complexity tends to be introduced through new feature needs, changes in size, and modification of the underlying infrastructure. The requirement to stay on top of emerging threats and changing assaults on the security front may also make systems more complicated. Additionally, the need to satisfy market expectations might push those responsible for developing and maintaining systems to take shortcuts and rack up technical debt.

Complexity often builds up unintentionally, yet this may result in tipping-point scenarios where a seemingly tiny change has significant repercussions for a system's security or stability. One well-known instance of a significant failure brought on by a little modification is a flaw that was introduced in the Debian GNU/Linux version of the OpenSSL library in 2006 and identified over two years later. An open-source programmer discovered that Valgrind, a common memory troubleshooting tool, was displaying warnings regarding RAM utilised before startup. The developer took off two lines of code in order to get rid of the warnings. Due to this, OpenSSL's pseudo-random number generator was regrettably only seeded with a process ID, which on Debian at the time was set to a value between 1 and 32,768 by default. Cryptographic keys might thus be broken with ease using brute force.

Even Google has had setbacks brought on by apparently unimportant modifications. For instance, a minor adjustment to a general logging library caused YouTube to be unavailable worldwide for more than an hour in October 2018. Both the author and the authorized code reviewer thought a modification to increase the level of event recording granularity was benign,

and it passed all tests. The update immediately caused YouTube servers to run out of memory and crash under production pressure, something the engineers were not fully aware of when they made the change. Cascading failures shut down the whole service as user traffic was diverted to other, still-healthy servers as a result of the failures.

### **Resilience**

Of course, a memory use issue shouldn't have led to a disruption in all services. Systems should be built to be robust in the face of unfavorable or unexpected conditions. Such situations are often brought on by unexpectedly large loads or component failures, from a reliability standpoint. You may accomplish resilience by either shedding some of the incoming load (processing fewer requests) or decreasing the processing cost for each request (processing more cheaply), since load is a function of the volume and average cost of requests to the system. System design should include redundancy and discrete failure zones to handle component failures so that you may reroute requests to minimise the effect of failures.

No matter how robust a system's individual parts may be, if it has enough complexity, it becomes difficult to prove that the whole thing is impervious to attack. Defence in depth techniques and specific failure domains might help you partially solve this issue. The use of several, perhaps redundant defence measures is known as defence in depth. Differentiated failure domains help to boost dependability by reducing the "blast radius" of failures. An adversary's capacity to use a compromised host or stolen credentials to move laterally, increase privilege, or have an impact on other system components is constrained by a sound system architecture. By separating permissions or limiting the scope of credentials, you may design several failure domains. For instance, the internal architecture of Google offers credentials that are specifically geographically scoped. These kinds of characteristics may restrict an attacker's ability to migrate laterally to servers in other areas after compromising a server in one location.

Another typical method for defence in depth is to use separate encryption layers for sensitive data. Discs, for instance, provide device-level encryption, but it's often a good idea to additionally encrypt the data at the application layer. In this approach, if an attacker has physical access to a storage device, even a defective implementation of an encryption algorithm in a drive controller won't be enough to jeopardise the confidentiality of protected data [7]–[9]. The examples provided thus far have all involved external attackers, but there are also internal risks that need to be taken into account. The two situations often don't vary much in practise, despite the fact that an insider may be more aware of possible misuse vectors than an outsider who is stealing an employee's credentials for the first time. Insider threats may be reduced by following the least privilege concept. A user must have the bare minimum of privileges necessary to carry out their duties at a certain moment, according to this rule. For instance, tools like sudo in Unix provide fine-grained restrictions that outline which users may execute which commands in which roles.

In order to make sure that sensitive actions are reviewed and authorized by certain groups of workers, Google additionally uses multi-party authorization. This multi-party system decreases the possibility of unintentional human mistake, a typical source of dependability problems, while also guarding against hostile insiders. The concepts of least privilege and multi-party authorization are not new; they have been used in a variety of non-computing contexts, such as bank vaults and nuclear missile silos.

### **From Design to Production**

Even a strong design should take security and dependability into account before being completely implemented as a production system. Through code reviews, beginning with the writing of the code itself, there are chances to identify possible security and reliability concerns, and even to stop the occurrence of whole classes of difficulties by making use of widely used frameworks and libraries. You may use testing to make sure a system works properly before deploying it, both in expected situations and in edge circumstances that often affect dependability and security. Testing plays a crucial role in ensuring that the system you've actually built adheres to your design intentions, whether you use load testing to understand how a system behaves under a barrage of queries, fuzzing to explore the behavior on potentially unexpected inputs, or specialized tests to make sure that cryptographic libraries aren't leaking information.

Last but not least, some methods for actually deploying code may reduce the risk of security and dependability. For instance, delayed rollouts and canaries can stop you from simultaneously damaging the system for all users. Similar to the last example, a deployment mechanism that only accepts code that has undergone a thorough review may assist to reduce the possibility of an insider putting a harmful binary into production [10].

### **Investigating Systems and Logging**

To yet, our attention has been on design tenets and implementation strategies that guard against both security and reliability issues. Unfortunately, achieving complete dependability or security is either unfeasible or too costly. You must prepare to identify and recover from failures because you must expect that preventative systems will malfunction. Effective logging is the cornerstone of readiness for failure and detection. In general, the more thorough and specific your logs are, the better but there are several exceptions to this rule. At a large enough scale, log volume presents a considerable expense, and efficient log analysis may become challenging. The YouTube example from earlier in this chapter demonstrates how dependability issues may also be brought on by logging. Security logs provide an extra problem since they normally shouldn't include sensitive data like login passwords or personally identifiable information (PII), lest they draw the attention of criminals.

### **Crisis Response**

Teams must function efficiently and rapidly during an emergency since errors might have instant repercussions. In the worst scenario, a situation may instantly ruin a company. For instance, in 2014 an attacker took control of the service's administrative tools and deleted all of its data, including all backups, putting the code-hosting service Code Spaces out of business in a couple of hours. Timely reactions in these circumstances depend on effective coordination and incident management.

It's easier to have a strategy in place before an issue arises since crisis response coordination is difficult. The time may have passed since the occurrence by the time you learn about it. Responders are, in any event, under stress, time constraints, and, at least initially, limited situational knowledge. The need to preserve state across teams and to pass over incident management at the borders of work shifts further complicates operations when a big organisation is involved and the problem calls for 24/7 response capabilities or cooperation across time zones.

Additionally, security incidents often include conflict between the desire to engage all necessary parties and the need—often prompted by legal or regulatory requirements—to limit information sharing to those who have a need-to-know basis. Furthermore, the original security breach could just be the beginning. The probe may extend beyond of the company's walls or include police enforcement.

Responders usually navigate lengthy stretches of time with minimal action when not under the strain of an active situation. Teams must maintain people's abilities and motivation throughout these periods while enhancing procedures and infrastructure in order to be ready for the next disaster. The Disaster Recovery Testing programme (DiRT) at Google frequently replicates a variety of internal system breakdowns and puts teams in difficult situations. Regular offensive security drills put our defences to the test and reveal brand-new weaknesses. Google uses IMAG even for minor accidents, which further motivates us to practice emergency procedures and tools on a regular basis.

### **Recovery**

Patching systems is often necessary to close a vulnerability in order to recover from a security breach. It seems sense that you would want that process to go as swiftly as feasible, utilising routinely used and hence mostly dependable techniques. The capacity to deploy updates fast, meanwhile, has a drawback: although it might help swiftly resolve vulnerabilities, it can also bring bugs or performance problems that cause significant harm. If the vulnerability is well-known or serious, there is more pressure to provide fixes promptly. It finally boils down to a risk assessment and a business decision as to whether to make remedies slowly—and so have greater confidence that there are no unintended side effects, but risk that the vulnerability will be exploited—or to do it rapidly. To repair a serious vulnerability, for instance, it can be appropriate to sacrifice some performance or increase resource utilization. The necessity for dependable recovery procedures that enable us to swiftly carry out critical modifications and upgrades without sacrificing dependability and that also notice possible issues before they create a large outage is highlighted by decisions like these. A trustworthy representation of each machine's current and intended state, as well as backstops to guarantee that state is never rolled back to an unreliable or dangerous version, are examples of requirements for a strong fleet recovery system.

### **CONCLUSION**

Security and dependability are intrinsic qualities of all information systems that are initially alluring to compromise in the name of speed, but expensive to rectify after the fact. As your systems develop and expand, this book seeks to assist you in addressing upcoming security and dependability concerns. Each organisation must comprehend the tasks and responsibilities that go along with engineering efforts in order to continue sustainable practices and create a culture of security and dependability. We aim to help you avoid paying a higher price later on by sharing our experiences and lessons gained so that you may embrace some of the guidelines outlined here early enough in the system lifespan. Regardless of the size or stage of your project, we produced this book with the intention that you would find it useful. While reading it, bear in mind the risk profile of your project; running a website for an animal sanctuary has a very different risk profile than running a stock market or a communication medium for dissidents. The enemies' many classifications and potential reasons are covered in length in the next chapter.

**REFERENCES**

- [1] L. J. Camp, "Designing for trust," in *Security and Privacy: Volume III*, 2016. doi: 10.4018/978-1-60566-264-0.ch026.
- [2] F. Gao, D. L. Chen, M. H. Weng, and R. Y. Yang, "Revealing development trends in blockchain-based 5g network technologies through patent analysis," *Sustain.*, 2021, doi: 10.3390/su13052548.
- [3] L. Bu, M. Mark, and M. Kinsy, "A short survey at the intersection of reliability and security in processor architecture designs," in *Proceedings of IEEE Computer Society Annual Symposium on VLSI, ISVLSI*, 2018. doi: 10.1109/ISVLSI.2018.00031.
- [4] A. Ahmadi, A. Esmaeel Nezhad, P. Siano, B. Hredzak, and S. Saha, "Information-Gap Decision Theory for Robust Security-Constrained Unit Commitment of Joint Renewable Energy and Gridable Vehicles," *IEEE Trans. Ind. Informatics*, 2020, doi: 10.1109/TII.2019.2908834.
- [5] P. Ercegovac, G. Stojić, M. Kopic, Ž. Stević, F. Sinani, and I. Tanackov, "Model for risk calculation and reliability comparison of level crossings," *Entropy*, 2021, doi: 10.3390/e23091230.
- [6] M. R. Ebling, E. De Lara, A. Wolman, and A. Gavrilovska, "The edge of the cloud," *IEEE Pervasive Computing*. 2013. doi: 10.1109/MPRV.2013.76.
- [7] M. Singh and S. Kim, "Branch based blockchain technology in intelligent vehicle," *Comput. Networks*, 2018, doi: 10.1016/j.comnet.2018.08.016.
- [8] I. A. Abbasi and E. E. Mustafa, "A Survey on Junction Selection based Routing Protocol for VANETs," *Int. J. Adv. Comput. Sci. Appl.*, 2021, doi: 10.14569/IJACSA.2021.0120121.
- [9] E. M. Biggs, J. M. A. Duncan, P. M. Atkinson, and J. Dash, "Plenty of water, not enough strategy," *Environ. Sci. Policy*, 2013, doi: 10.1016/j.envsci.2013.07.004.
- [10] A. M. Mansoor, A. M. Sarea, and A. Q. Md Sabri, "Influence of intersections on the performance of position-based routing protocols for VANETs," *World J. Entrep. Manag. Sustain. Dev.*, 2018, doi: 10.1108/wjemdsd-08-2017-0056.

## A BRIEF INTRODUCTION ON UNDERSTANDING ADVERSARIES

**Ms. Appaji Gowda Shwetha\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:shwetha.a@presidencyuniversity.in

---

### ABSTRACT:

*Building successful defence measures in the field of cybersecurity requires a thorough knowledge of the adversary. The term "adversaries," sometimes known as "threat actors," refers to a broad spectrum of organisations, including hackers, cybercriminal groups, nation-state actors, and insider threats. Organisations may improve their capacity to identify, stop, and react to cyber-attacks by acquiring insights into the motives, methods, approaches, and targets of these threats. The significance of comprehending enemies in the context of cybersecurity is explored in this abstract. It emphasizes the need for businesses to go above conventional security measures and take a proactive stance that emphasizes threat intelligence, adversary profiling, and constant monitoring.*

**KEYWORDS:** *Cybersecurity, Intelligence Gathering, Risk Management Vulnerability Researchers.*

---

### INTRODUCTION

The existence of adversaries is a key issue that organisations must address in the area of cybersecurity and risk management. These adversaries, also known as threat actors, may be anything from lone hackers to well-organized gangs to even state-sponsored organisations. A key component of creating strong defences and reducing possible dangers is understanding the intentions, strategies, and methods used by these adversaries. Obtaining and examining data on an adversary's capabilities, objectives, and possible attack routes is necessary for understanding them. Organisations may use this information to determine the precise dangers they face, the possible consequences of an assault, and the weaknesses that adversaries can use against them. Organisations may create specialized defences and manage resources efficiently by developing a greater knowledge of their enemies.

Grasp enemies requires a thorough grasp of threat intelligence. Data from several sources, including open-source intelligence, dark web surveillance, and information-sharing platforms, are gathered and analyzed in this process. Organisations may keep updated about new threats, developing attack methods, and the strategies used by various adversaries by using threat intelligence.

Another important aspect of knowing enemies is adversary profiling. It entails the creation of profiles that describe the traits, motives, and capabilities of recognized threat actors. Through the use of these profiles, organisations may deploy targeted defences and responses by anticipating the actions and strategies that potential adversaries may take.



To keep a current awareness of opponents, continuous monitoring is necessary. Organisations should set up reliable monitoring systems and use cutting-edge analytics approaches to find unusual activity and possible signs of compromise. Organisations may proactively detect possible threats and take necessary action to limit risks by continually monitoring their systems, networks, and data. Understanding opponents requires more than just technical knowledge; it also requires taking into account the larger environment in which they operate. This covers the influence of developing technology, as well as geopolitical and economic considerations. Understanding these outside factors enables organisations to more accurately gauge the capabilities and intentions of their enemies and adjust their defences as necessary. It is crucial in the realm of cybersecurity to comprehend opponents. Organisations may improve their security posture and successfully safeguard their systems and data by getting insights into their adversaries' motives, strategies, and goals. Organisations can keep ahead of developing threats, foresee attack pathways, and apply tailored defences via threat information, adversary profiling, and continuous monitoring. Organisations may reduce risks in the dynamic environment of cybersecurity by making educated choices, allocating resources wisely, and acting quickly in response to threats.

Organisations must go beyond conventional security measures to understand the motivations of prospective attackers in order to genuinely secure sensitive information and valuable assets. Understanding their intentions, talents, and techniques for exploiting weaknesses is part of this process. Organisations may strengthen their security posture by better anticipating attacks and acting proactively by knowing their enemies. In the digital world, adversaries might have a variety of goals. Some engage in illegal acts including data theft, fraud, and ransomware attacks in order to make money. Other people could try to damage vital infrastructure, further their own political or ideological objectives, or carry out espionage. Organisations may priorities their security tactics and customize their defences by understanding the precise motives of various threat actors [1]–[3].

Furthermore, it's critical to understand the strategies and methods that attackers use in order to find possible security holes in networks and systems. Adversaries constantly modify and improve their strategies, using both advanced technology and social engineering to get beyond security precautions. Organisations may successfully combat assaults by using strong countermeasures, such as intrusion detection systems, enhanced monitoring, and security awareness training, by remaining informed of these techniques. Analysing an adversary's capabilities is essential to understanding them. Using sophisticated malware, zero-day vulnerabilities, or advanced persistent threats (APTs), certain adversaries may have highly developed technical capabilities. Others may use tricks of deceit and social engineering to prey on flaws in people. Organisations may improve their incident response skills, create efficient incident management processes, and adopt suitable security policies to limit risks by getting knowledge into the capabilities of their opponents.

Understanding enemies also requires coordination and the exchange of threat knowledge. Information on the most recent threat patterns, signs of compromise, and attack tactics is often exchanged between industry groups, cybersecurity organisations, and government authorities. Organisations may acquire information and insights by engaging in such cooperative projects, allowing them to better foresee and fight against new dangers.

## DISCUSSION

For a system to be resilient and survive a broad range of disasters, its opponents must be understood. In the context of dependability, opponents often act with good intentions and adopt an abstract shape. They may manifest as common technology malfunctions or instances of disproportionately high user interest (referred to as "success disasters"). They might also be fishing boats that unintentionally cut fiber-optic connections beneath the sea or configuration modifications that make systems act unexpectedly. In contrast, security context adversaries are human; their activities are planned to negatively impact the target system. Despite these divergent objectives and approaches, understanding reliability and security threats is crucial for developing resilient system designs and implementations. Without this information, it would be very difficult to predict the behavior of a Wily Hacker or a Curious Cat.

In this chapter, we go in-depth on security adversaries to assist experts from a variety of industries in acquiring an adversarial perspective. It may be alluring to see security threats through the prism of common stereotypes: assailants in pitch-black basements with catchy monikers and maybe dubious behavior. Even while there are surely colorful characters like these, anybody with enough time, information, or resources may compromise a system's security. Anyone with physical access to a computer or mobile device may buy software that lets them to take control of the device for a minimal charge. Governments often purchase or create software to jeopardise the targets' systems. In order to better understand how systems function, researchers often test their safety features. As a result, we urge you to have an impartial viewpoint on the source of system attacks.

There are never two identical assaults or assailants. For a discussion of the cultural nuances of dealing with foes, we suggest reading Chapter 21. Even for experienced security professionals, predicting impending security calamities is mostly an exercise in guesswork. In the parts that follow, we give three frameworks for comprehending attackers that we have found useful over the years. These frameworks examine probable reasons why individuals would attack systems, some typical attacker characteristics, and how to think about attackers' techniques. Within the three frameworks, we also provide instructive (and hopefully interesting) examples.

### Attacker Motivations

Security threats are mostly human (for the time being at least). As a result, we may see the motivation for assaults from the perspective of those who do them. By doing this, we could be better able to comprehend how we ought to react in both a proactive (during system design) and reactive (during accidents) manner.

An attacker may be a criminal actor, a government snoop, and a financially motivated vulnerability researcher all at once! For instance, the US Department of Justice indicted Park Jin Hyok in June 2018, a North Korean citizen accused of engaging in a number of activities on behalf of his government, such as developing the notorious WannaCry Ransomware in 2017 (used for financial gain), hacking Sony Pictures in 2014 (intended to force Sony into withholding the release of a contentious film, ultimately harming the company's infrastructure), and hacking electric utilities (predatory). Researchers have also seen government attackers steal digital currency from video games for their own benefit using the same malware used in nation-state operations.

These many motives should be considered while creating systems. Take into account a business that manages money transfers for its clients. We can build the system more securely if we know the potential motivations of an attacker. The actions of a group of North Korean government attackers (including Park) who allegedly tried to steal millions of dollars by hacking banking systems and using the SWIFT transaction system to transfer money out of customer accounts provide a good illustration of potential motivations in this case.

### **Defending Attacker Profiles**

By considering the individuals themselves, including who they are, whether they carry out assaults for themselves or someone else, and their general interests, we may better comprehend the motives of attackers. This section includes some profiles of attackers, explanations of how they relate to system designers, and advice on how to guard your systems against these kinds of attackers. Remember that no two assaults or attackers are alike despite the generalizations we've used for the purpose of conciseness. This data is intended to be indicative rather than conclusive [4]–[6].

### **Hobbyists**

The original computer hackers were amateur technologists with a curiosity for how things worked. These "hackers" found bugs that the original system designers had missed when disassembling computers or troubleshooting their software. In general, hobbyists are driven by their desire to learn; they hack for amusement and may work with engineers to include robustness into a system. The majority of the time, enthusiasts follow personal principles that prohibit hurting systems and refrain from engaging in illegal activity. You may increase the security of your systems by using knowledge of the problem-solving processes used by these hackers.

### **Vulnerability Researchers**

Researchers that specialize on security issues do so professionally. As full-time workers, part-time independent contractors, or even accidentally as regular users who discover vulnerabilities, they take pleasure in identifying security issues. The Vulnerability Reward Programmes, commonly referred to as bug bounties, are popular among researchers. Researchers that focus on system vulnerabilities are often driven to improve systems and are valuable partners for businesses looking to safeguard their infrastructure. The way vulnerabilities are found, disclosed, corrected, and discussed between system owners and researchers is often expected to follow a set of predictable disclosure rules. These guidelines prevent researchers from improperly obtaining data, harming people, or breaching the law. Operating outside of these rules often renders the potential of receiving a reward illegitimate and may constitute illegal activity.

Red Teams and penetration testers, who may be specifically employed for these exercises, attack targets with the consent of the system owner. They follow a set of ethical rules and, like researchers, seek for methods to undermine system security with an emphasis on enhancing security.

### **Governments and Law Enforcement**

Security professionals may be hired by government organisations (such as law enforcement and intelligence agencies) to acquire information, combat domestic crime, conduct economic

espionage, or support military operations. For these reasons, the majority of national governments have now made investments in developing security expertise. Governments may sometimes resort to gifted recent graduates, ex-offenders who have changed their ways and served time in prison, or well-known figures in the security sector. Although we are unable to discuss these attackers in great detail here, we do provide a few instances of some of their more typical behaviors.

### **Intelligence gathering**

The government function that employs those skilled in system hacking, intelligence collection, is perhaps the one that is mentioned in the public eye the most. With the development of the internet over the last several decades, classic espionage methods such as human intelligence (HUMINT) and signals intelligence (SIGINT) have been modernized. In one well-known incident from 2011, the security firm RSA was breached by an adversary who many analysts believe to be a part of China's spy network. For the purpose of stealing the cryptographic seeds for RSA's well-known two-factor authentication tokens, the attackers exploited RSA. Once they had these seeds, the attackers were able to create one-time authentication credentials without the requirement for actual tokens, allowing them to access Lockheed Martin's systems, a defence contractor that produces equipment for the US military. In the past, breaking into a corporation like Lockheed would have been carried out on-site by human operatives—for instance, by paying off an employee or placing a spy inside the company. However, the emergence of systems penetration has given attackers the ability to utilise more advanced electrical approaches to gain secrets in novel ways.

### **Military purposes**

System hacking by governments for military objectives is known as cyberwarfare or information warfare by experts. Imagine a government planning an invasion of another nation. Could they possibly thwart the air defences of the target and make them fail to detect an approaching air force? Could they turn off their banking, water, and electricity systems? As an alternative, consider a scenario in which a government aims to stop a different nation from developing or acquiring weapons. Could they sabotage their growth covertly and remotely? According to reports, this scenario took place in Iran in the late 2000s when hackers secretly installed a modularized piece of software on the management systems of centrifuges used for uranium enrichment. Researchers have given this operation, known as Stuxnet, the goal of stopping Iran's nuclear programme by destroying its centrifuges.

### **Policing domestic activity**

Governments may potentially hack into domestic activity monitoring systems. In a recent instance, the cybersecurity contractor NSO Group offered software to multiple countries that permitted covert remote monitoring of mobile phone conversations, allowing private surveillance of individual communications without the target's awareness. This programme allegedly had the purpose of tracking criminals and terrorists two rather uncontroversial targets. The ethics of governments using these capabilities against their own people is a hotly debated topic, especially in nations without strong legal frameworks and adequate oversight. Unfortunately, some of NSO Group's government customers have also used the software to listen in on journalists and activists, which has in some cases resulted in harassment, arrest, and even death<sup>3</sup>.

### **Protecting your systems from nation-state actors**

System developers need to carefully assess if they may be a nation-state actor's target. To that purpose, you must comprehend the organizational actions that can appeal to these players. Think about a technological business that manufactures and distributes microprocessors to the military. It's feasible that other nations would be interested in obtaining such chips as well and would even resort to electronically stealing their blueprints [7]–[9].

Additionally, your service can contain information that the government needs but finds challenging to get. In general, police enforcement and intelligence organisations value private conversations, location information, and comparable categories of sensitive personal data. It is now generally accepted that the sophisticated targeted assault from China on Google's corporate infrastructure in January 2010 dubbed "Operation Aurora" by researchers was intended to acquire long-term access to Gmail accounts. Private communications in particular may increase the possibility that an intelligence or law enforcement organisation will be interested in your systems if you save the personal information of your clients.

You could unknowingly become a target from time to time. Operation Aurora impacted at least 20 victims from a range of the financial, technology, media, and chemical industries in addition to big IT firms. These organisations ranged in size from big to tiny, and many did not believe they were vulnerable to assault by a nation-state. Take into account, for instance, an app that attempts to provide athletes data monitoring statistics, including the locations where they run or bike. Would an intelligence agency find this data to be a desirable target? When they discovered that US personnel were using the service to monitor their workouts, the whereabouts of covert military sites in Syria were made public on a heat map made by the fitness tracking startup Strava in 2018, analysts wondered just this.

The use of large resources by governments to get access to data is a common practice that system designers should be aware of. The amount of resources needed to mount a defence against a government that wants access to your data may be substantially more than what your company can allocate to putting security measures in place. We advise businesses to adopt a long-term approach to creating security defences by investing early in safeguarding their most valuable assets and by maintaining a strict programme that may add additional layers of security as time goes on. An ideal result would be to have an enemy spend a lot of resources on you, increasing their chance of being discovered, so that their actions may be made known to other potential victims and government authorities.

### **Activists**

Hactivism is the practice of requesting social change using technology. This word is used indiscriminately to refer to a broad range of online political actions, from the subversion of government monitoring to the deliberate disruption of systems. Websites have been known to be defaced by hackers, who alter the original material with political commentary. One instance from 2015 saw the Syrian Electronic Army, a group of hostile actors working in favor of the Bashar al-Assad government, seizing control of a content distribution network (CDN) that handled online traffic for [www.army.mil](http://www.army.mil). The attackers were then able to incorporate a message in favor of Assad, which website users later saw. Attacks of this kind may be very humiliating for website owners and damage visitors' faith in the platform.

More damaging hacktivist assaults could exist. For instance, using denial-of-service operations in November 2012, the multinational, decentralized hacktivist collective Anonymous5 brought multiple Israeli websites down. Anyone accessing the impacted websites therefore encountered a delay in service or an error. These kinds of distributed denial-of-service assaults bombard the target with traffic from tens of thousands of infected devices dispersed throughout the globe. This kind of capacity is often offered for sale online by brokers of these so-called botnets, making assaults prevalent and simple to carry out. Attackers may even threaten to completely damage or disrupt systems, which has led some experts to classify them as cyber terrorists.

Hacktivism is often open about their actions and frequently claim credit in public, in contrast to other sorts of attackers. This may take several forms, such as publishing on social media or dismantling institutions. Even the attackers themselves may not be very skilled technically. This might make it challenging to anticipate or protect against hacktivism.

### **Protecting your systems from hacktivists**

We advise you to consider if your company or project is engaged in contentious issues that might interest protestors. Does your website, for instance, let users to upload their own material such as blogs or videos? Does your initiative touch on a topic with political overtones, such as animal rights? Are any of your goods, such a messaging service, used by activists? If the answer to any of these questions is "yes," you may want to think about installing extremely strong, layered security measures that make sure your systems are patched against flaws, resistant to DoS assaults, and that backups can swiftly restore a system and its data.

### **Criminal Actors**

Attack methods are used to crimes that closely mirror those committed by their non-digital counterparts, such as committing identity fraud, stealing money, and blackmail. Criminal actors are skilled in a variety of technical areas. Some people may be knowledgeable and create their own tools. Others may use tools that other people have built, depending on their simple, click-to-attack interfaces to buy or borrow them. In spite of being at the lowest level of complexity, social engineering, which involves persuading a victim into helping you in the assault, is quite successful. Most criminals simply need a small amount of time, a computer, and some money to get started.

It would be hard to provide a comprehensive list of the many illegal behaviors that take place online, but we do give a few sample instances instead. Consider the scenario if you wished to anticipate merger and acquisition activity so that you could schedule certain stock transactions appropriately. This same plan was devised by three criminals in China in 2014–2015, who profited a few million dollars by stealing confidential information from unwary legal firms.

Attackers have learned over the previous ten years that when a victim's sensitive data is at risk, they will accept payment. Until the victim pays the attacker, ransomware is malware that keeps a machine or its data hostage (often by encrypting it). By taking advantage of flaws, bundling ransomware with safe software, or persuading the user into downloading it themselves, attackers often infect target computers with this malware (which is frequently packaged and marketed to attackers as a toolkit) [10].

Criminal behavior may not necessarily take the form of overt efforts to commit theft. The purpose of stalkerware, which can be purchased for as little as \$20, is to acquire data on another

person without that person's awareness. Either the victim is tricked into downloading the malicious software or an attacker with access to the device directly installs it on the victim's PC or mobile device. The programme may record audio and video after it is installed. This kind of trust exploitation may be very successful since stalkerware is often employed by somebody close to the victim, such a spouse.

Criminals don't always work for themselves. For their own ends, businesses, legal firms, political campaigns, cartels, gangs, and other groups use bad actors. For instance, a Colombian assailant said he was recruited to help a candidate in the 2012 Mexican presidential election as well as other elections held around Latin America by taking material from the opposition and disseminating rumours. In a startling example from Liberia, a Cellcom employee allegedly hired an attacker to weaken the network of Lonestar, a competitor cellular service provider. Due to the assaults' interference with LoNestar's capacity to provide for its clients, the firm suffered huge financial losses.

### **Protecting your systems from criminal actors**

Remember that criminal actors tend to gravitate towards the simplest approach to achieve their aims with the least upfront cost and effort when developing systems to be robust against them. If you can build a robust enough system, they could decide to concentrate on a different victim. Take this into account when deciding which systems to target and how to make their assaults costly. An effective example of how to gradually raise the cost of assaults is the development of Completely Automated Public Turing test (CAPTCHA) systems. The purpose of CAPTCHAs is to distinguish between human users and automated bots when they interact with a website, such as when logging in. Knowing if a user is human may be a key indicator since bots are often an indication of criminal activities. Early CAPTCHA systems required humans to authenticate slightly crooked characters or numbers that computer programmes had trouble reading. Implementers of CAPTCHA started using distortion images and object recognition as the bots evolved. These strategies were designed to gradually raise the expense of breaking CAPTCHAs.

### **Automation and Artificial Intelligence**

The US Defence Advanced Research Projects Agency (DARPA) announced the Cyber Grand Challenge competition in 2015 with the goal of creating a cyber-reasoning system that could operate autonomously and self-learn to identify software flaws, devise methods to exploit them, and then develop countermeasures. Seven teams competed in a live "final event" while watching their completely autonomous reasoning systems battle it out in a spacious ballroom. The winning team was successful in creating such a self-learning system. The Cyber Grand Challenge's accomplishment raises the possibility that at least some assaults in the future may be carried out without direct human involvement. whether fully sentient robots were to exist, some scientists and ethicists wonder whether they may be able to teach one another how to fight. We anticipate that the concept of autonomous assault platforms will lead to a need for more automated defences, which will be a significant topic of study for system designers in the future.

### **CONCLUSION**

Every security breach may be linked to a motivated individual. To assist you in determining who would wish to target your services and why, we've covered some typical attacker characteristics. This will help you priorities your defences. Determine any potential targets. What do you have to

offer? Who purchases your goods and services? Could the behavior of your users inspire attackers? How do your defensive resources stack up against the attacking capabilities of prospective rivals? The knowledge in the next chapters of this book may assist make you a more costly target even when you are dealing with a well-funded opponent, thereby eliminating the financial motive for an assault. Don't ignore the smaller, less noticeable opponent; obscurity, location, plenty of time, and the difficulties of prosecution may all be benefits to an assailant, enabling them to harm you disproportionately severely. Because all organisations might potentially face insider threats, both malevolent and not, it is important to think about your insider risk. Insiders have privileged access that enables them to do serious harm.

## REFERENCES

- [1] M. McFate, "The Military Utility of Understanding Adversary Culture," *JFQ Jt. Force Q.*, 2005.
- [2] R. Derbyshire, B. Green, and D. Hutchison, "'Talking a different Language': Anticipating adversary attack cost for cyber risk assessment," *Comput. Secur.*, 2021, doi: 10.1016/j.cose.2020.102163.
- [3] B. Haykel, "ISIS and al-Qaeda—What Are They Thinking? Understanding the Adversary," *Ann. Am. Acad. Pol. Soc. Sci.*, 2016, doi: 10.1177/0002716216672649.
- [4] M. Carrier, R. Mertens, and C. Reinhardt, *Narratives and comparisons: Adversaries or allies in understanding science?* 2021. doi: 10.14361/9783839454152.
- [5] Z. Deng, H. He, J. Huang, and W. J. Su, "Towards understanding the dynamics of the first-order adversaries," in *37th International Conference on Machine Learning, ICML 2020*, 2020.
- [6] W. Young, "Understanding and Diagnosing Adversary System Behavior in 4th Generation Warfare: A Soft Approach to EBO Mission Analysis," *Proceedings of the 23rd International Conference of the System Dynamics Society*. 2005.
- [7] K. D. Chowdhury and D. Elliott, "Understanding the effect of textual adversaries in multimodal machine translation," in *LANTERN@EMNLP-IJCNLP 2019 - Beyond Vision and LANGUAGE: inTEgrating Real-World kNowledge, Proceedings*, 2019. doi: 10.18653/v1/d19-6406.
- [8] C. Martini, "What 'Evidence' in Evidence-Based Medicine?," *Topoi*, 2021, doi: 10.1007/s11245-020-09703-4.
- [9] UK Ministry of Defence, "Understanding and Intelligence Support to Joint Operations (JDP 2-00)," *Jt. Doctrin. Publ.*, 2011.
- [10] Criag Smith, "Next Generation Red Teaming," *Netw. Secur.*, 2016, doi: 10.1016/s1353-4858(16)30035-6.



## A CASE STUDY ON SAFE PROXIES

**Mr. Bhavan Kumar\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:bhavankumar.m@presidencyuniversity.in

---

### **ABSTRACT:**

*In this case study, the use of safe proxies is examined within the framework of site reliability engineering (SRE). In order to guarantee the dependability, security, and performance of software systems, safe proxies are an essential element. They provide as a bridge between users and services, managing traffic, distributing load, and enforcing security regulations. The paper starts out by examining the difficulties that organisations have when trying to manage and secure client-server communication. It emphasizes the need of safe proxies in dealing with these issues and gives a general overview of how they help to increase system dependability and reduce hazards. The case study then explores how safe proxies are implemented inside an SRE architecture. It covers the variables to take into account when choosing a suitable proxy solution, including scalability, performance, security, and compatibility with current infrastructure.*

**KEYWORDS:** *Compatibility, Performance, Scalability, Security.*

---

### **INTRODUCTION**

The subject of site reliability engineering (SRE) is concerned with preserving and enhancing the dependability, availability, and performance of software systems. The use of proxies in managing and safeguarding system communications has grown in significance as part of SRE practices. This case study examines the use of safe proxies in an SRE setting and demonstrates their advantages in boosting security and dependability. Safe proxies serve as a barrier between client requests and backend services, adding an extra measure of security and management. They are essential for traffic routing and load balancing, performance optimization, and risk mitigation. Organisations may maintain the efficient operation of their systems while protecting against nefarious activity and interruptions by properly setting and controlling these proxies. The research also looks at safe proxy deployment methodologies and architectural design. It examines several deployment models, including sidecar, forward, and reverse proxies, and considers the benefits and drawbacks of each. Additionally, it covers issues including traffic routing systems, load balancing algorithms, and service discovery. The case study then explores the safe proxies' security features. It talks about how crucial authentication, access control, and encryption are to safeguarding private information and preventing unauthorized access. It also examines methods for preventing typical security concerns including Distributed Denial of Service (DDoS) assaults.

The paper also emphasizes the monitoring and observability features of safe proxies. In order to properly troubleshoot problems, detect bottlenecks, and get insights about proxy performance, it

---

emphasizes the necessity for robust logging, metering, and tracing capabilities. In the case study's conclusion, organisations who have included safe proxies into their SRE practices are given real-world examples and success stories. It emphasizes the advantages they have attained, such as increased security, greater dependability, and simplified operations. Overall, this case study offers a thorough investigation of safe proxies in relation to SRE. For businesses looking to use secure proxies to raise the dependability and security of their software systems, it provides insightful advice and useful tips.

Improved dependability is among the safe proxies in SRE's main advantages. Safe proxies divide incoming traffic across many backend services using load balancing methods, avoiding any one service from being overloaded. This helps prevent service degradation or outages brought on by too many requests or sudden increases in traffic. Proxy servers may also assess the status of backend services, automatically diverting traffic from those that are unhealthy or failing, and keeping the system available.

Safe proxies operate as a security barrier between users and the backend services, protecting them from direct exposure to dangers from the outside world. Only authorized customers will be able to access the services thanks to the enforcement of access rules, authentication, and authorization processes through proxies. Incoming requests may also be examined and filtered, allowing for the detection and prevention of harmful activity like distributed denial-of-service (DDoS) assaults and efforts to exploit security flaws. Proxies may further safeguard sensitive data in transit by using secure protocols and encryption [1]–[3].

The adaptability of safe proxies in handling system upgrades and changes is another important benefit. Proxies may dynamically direct traffic to new or improved backend services, enabling smooth deployments and preventing outages. In order to enhance speed and lower network latency, they may also include request and response modification capabilities, allowing transformations, caching, and content optimization. Safe proxy implementation in an SRE setting requires meticulous design, setup, and monitoring. To create a successful proxy architecture, SRE teams must take into account variables including traffic patterns, scalability needs, security standards, and performance goals. To spot abnormalities or possible security breaches, proxy performance, traffic patterns, and security logs must be regularly monitored.

## **DISCUSSION**

### **Safe Proxies in Production Environments**

Proxy servers often provide a mechanism to handle new security and reliability needs without significantly altering the existing systems. You may simply use a proxy to redirect connections that would have otherwise gone straight to the system rather than making changes to an existing system. Controls to satisfy your updated security and dependability needs may also be included in the proxy. In this case study, we look at a set of secure proxies that Google employs to prevent privileged administrators from deliberately or unintentionally disrupting our production environment.

A system known as a "safe proxy" enables authorized individuals to access or change the status of real servers, virtual machines, or specific applications. In order to evaluate, authorize, and execute dangerous instructions without first establishing an SSH connection to a system, Google uses safe proxies. We may rate-limit system restarts or offer granular access to diagnose

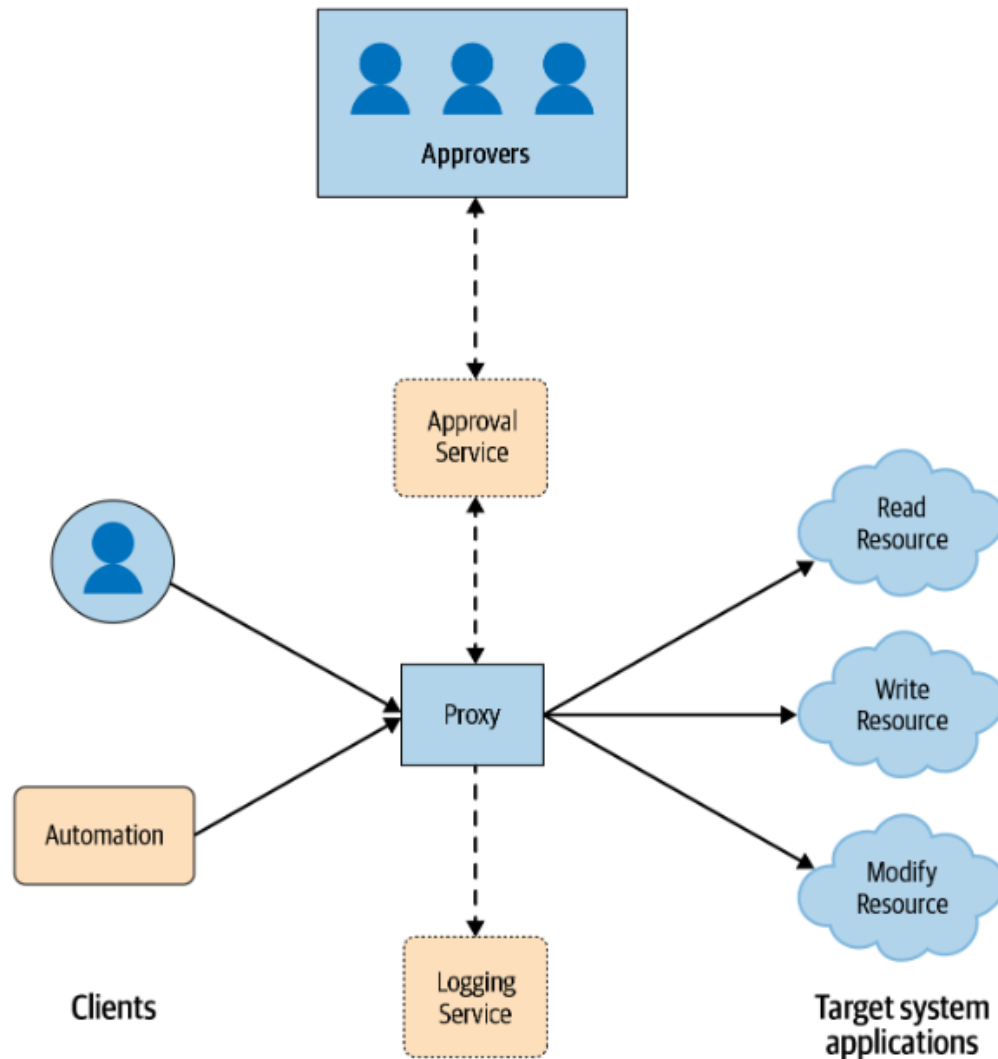
problems using these proxies. Safe proxies act as a single point of access across networks and are essential tools that let us:

1. Examine every fleet action.
2. Limit who has access to what
3. Prevent large-scale manufacturing from human error

As part of the Zero Touch Prod effort at Google, all production changes must be automated (rather being done by a person), prevalidated by computer software, or started by an audited break glass mechanism. Among the technologies we use to carry out these principles are safe proxies. We estimate that Zero Touch Prod might have avoided or reduced around 13% of the outages that Google investigated.

Clients communicate with the proxy in the safe proxy paradigm shown in Figure 1 rather than the target system directly. At Google, we impose this behavior by configuring the target system to only accept calls from the proxy. Through access control lists (ACLs), this setting determines which application-layer remote procedure calls (RPCs) may be conducted by which client roles. The proxy delivers the request to be processed through the RPC to the target systems after confirming the access permissions. Typically, an application-layer programme on each target system gets the request and runs it directly on the machine. All requests and instructions sent by the systems with which the proxy communicates are logged.

- a. Using proxies to administer systems has a number of advantages, whether the client is a person, an automated system, or both.
- b. A central location where we decide whether to provide access to requests that interact with sensitive data in order to impose multi-party authorization (MPA).
- c. Using administrative use auditing, we may monitor who executed a certain request at what time.
- d. Rate limitation, which progressively implements changes like a system restart, may allow us to reduce the explosion radius of an error.
- e. Compatibility with closed-source third-party target systems, where we govern the behavior of components (that we cannot edit) using proxy functionality.
- f. Integrating continuous development, we strengthen the primary proxy point's security and dependability.



**Figure 1: Illustrate the Safe proxy model.**

Proxies have various drawbacks and possible pitfalls:

- i. Increased maintenance and operating overhead costs.
- ii. If the system or one of its dependents fails, there is only one point of failure. We counteract this by running numerous instances to boost redundancy. We ensure that all of our system's dependents have an appropriate service level agreement (SLA) and that each of the dependencies' teams has a documented emergency contact.
- iii. A policy setting for access control that may itself be a cause of mistakes. We help consumers make the correct decisions by giving templates or automatically creating secure settings by default. We use the design methods discussed in Part II for constructing such templates or automation.
- iv. A central machine that might be taken over by an opponent. The aforementioned policy setup requires the system to relay the client's identification and perform any actions on


the client's behalf. Because no request is performed under a proxy role, the proxy itself does not provide high privileges.

- v. Users may be resistant to change because they want to link directly to production systems. To lessen the friction caused by the proxy, we collaborate closely with engineers to ensure that they may access the systems during crises through a breakglass mechanism.

Because the safe proxy's primary use case is to offer security and reliability features related to access control, the proxy's interfaces should utilise the same external APIs as the destination system. As a consequence, the proxy has no effect on the user experience. Assuming the safe proxy is transparent, it may simply forward traffic after doing some validation and recording pre- and postprocessing. The next section goes through one particular instance of a secure proxy that we utilise at Google.

### Proxy Google Tool

The bulk of administrative tasks are carried out by Googlers utilising command-line interface (CLI) tools. Some of these tools are potentially harmful; for example, some tools have the ability to shut down a server. If such a tool provides an inaccurate scope selection, a command-line invocation may cause numerous service frontends to cease, resulting in an outage. It would be difficult and costly to monitor every CLI tool, guarantee that it conducts centralized logging, and ensure that sensitive activities have additional safeguards. To overcome this problem, Google developed a Tool Proxy, which is a binary that exposes a generic RPC function that internally runs the provided command line through fork and exec. All invocations are governed by a policy, are recorded for auditing, and may require MPA[4]–[6]. Using the Tool Proxy meets one of Zero Touch Prod's key goals: making production safer by preventing people from directly accessing production. Engineers cannot execute arbitrary commands directly on servers; instead, they must use the Tool Proxy[7]–[9].



**Figure 2: Tool Proxy usage workflow.**

The Tool Proxy necessitates a little modification to the development workflow: engineers must prefix their commands with `tool-proxy-cli --proxy_address`. We updated the server to allow only administrative activities to `admin-proxy` and to disallow any direct connections outside of breakglass scenarios to prevent privileged users from circumventing the proxy[10].

### CONCLUSION

One method for adding logging and multi-party authorization to a system is to use safe proxies. Proxies may therefore assist in making your systems more safe and trustworthy. This strategy may be a low-cost solution for an existing system, but it will be considerably more robust when

combined with other design ideas. If you're beginning a new project, your system architecture should ideally be built utilising frameworks that interface with logging and access control modules. A breakglass mechanism is one that allows engineers to swiftly address outages by circumventing restrictions. MPA needs an extra user to authorize an action before it can take place. MPA stands for Multi-Party Authorization.

## REFERENCES

- [1] A. W. Hjalsted, A. Laurent, M. M. Andersen, K. H. Olsen, M. Ryberg, and M. Hauschild, "Sharing the safe operating space: Exploring ethical allocation principles to operationalize the planetary boundaries and assess absolute sustainability at individual and industrial sector levels," *J. Ind. Ecol.*, 2021, doi: 10.1111/jiec.13050.
- [2] L. K. Dwivedi, K. Banerjee, N. Jain, M. Ranjan, and P. Dixit, "Child health and unhealthy sanitary practices in India: Evidence from Recent Round of National Family Health Survey-IV," *SSM - Popul. Heal.*, 2019, doi: 10.1016/j.ssmph.2018.10.013.
- [3] A. Nielsen *et al.*, "Trial protocol: A parallel group, individually randomized clinical trial to evaluate the effect of a mobile phone application to improve sexual health among youth in Stockholm County," *BMC Public Health*, 2018, doi: 10.1186/s12889-018-5110-9.
- [4] B. S. Husebo, H. L. Heintz, L. I. Berge, P. Owoyemi, A. T. Rahman, and I. V. Vahia, "Sensing technology to facilitate behavioral and psychological symptoms and to monitor treatment response in people with dementia: A systematic review," *Frontiers in Pharmacology*. 2020. doi: 10.3389/fphar.2019.01699.
- [5] L. Frost, "Water Technology and the Urban Environment: Water, Sewerage, and Disease in San Francisco and Melbourne before 1920," *J. Urban Hist.*, 2020, doi: 10.1177/0096144217692988.
- [6] J. Neille, "A Qualitative inquiry into the ways in which space and place influence the lived experiences of adults with disabilities in rural south Africa," *Rural Remote Health*, 2021, doi: 10.22605/RRH6241.
- [7] G. G. G. Donders, K. Ruban, G. Bellen, and S. Grinceviciene, "Pharmacotherapy for the treatment of vaginal atrophy," *Expert Opinion on Pharmacotherapy*. 2019. doi: 10.1080/14656566.2019.1574752.
- [8] C. Chen, L. Liu, and N. Zhao, "Fear Sentiment, Uncertainty, and Bitcoin Price Dynamics: The Case of COVID-19," *Emerg. Mark. Financ. Trade*, 2020, doi: 10.1080/1540496X.2020.1787150.
- [9] A. Al Naimi, P. Moore, D. Brüggmann, L. Krysa, F. Louwen, and F. Bahlmann, "Ectopic pregnancy: a single-center experience over ten years," *Reprod. Biol. Endocrinol.*, 2021, doi: 10.1186/s12958-021-00761-w.
- [10] S. C. L. Watson, N. J. Beaumont, S. Widdicombe, and D. M. Paterson, "Comparing the network structure and resilience of two benthic estuarine systems following the implementation of nutrient mitigation actions," *Estuar. Coast. Shelf Sci.*, 2020, doi: 10.1016/j.ecss.2018.12.016.

## A BRIEF INTRODUCTION ON DESIGN TRADEOFFS

**Dr. Nakul Ramanna Sanjeevaiah\***

\*Associate Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA,  
Email Id:nakul@presidencyuniversity.in

---

### ABSTRACT:

*SRE has developed as a critical discipline for assuring the reliability, availability, and performance of contemporary software systems. SRE principles emphasize the need of engineering practices that balance operational duties with new feature development. Making design choices to meet these goals while successfully managing limited resources is one of the most difficult tasks in SRE. This paper provides a thorough examination of the design choices confronting Site Reliability Engineers. It delves into many aspects where compromises are typical, including as scalability, fault tolerance, latency, cost optimization, and user experience. To give practical insights, the examination includes both theoretical frameworks and real-world case studies. The research looks at the tradeoffs that must be made while selecting architectural patterns, infrastructure technologies, deployment methodologies, and monitoring and alerting systems. It investigates how these choices affect important performance measures such as system uptime, response time, and customer happiness. Furthermore, it investigates tradeoffs in capacity planning, load balancing, caching methods, and data persistence systems, taking into account the specific needs of various applications and services.*

**KEYWORDS:** *Design Tradeoffs, Scalability, Fault Tolerance, Latency, Cost Optimization, User Experience.*

---

### INTRODUCTION

Security and dependability criteria are sometimes challenging to balance with project feature and cost constraints. This chapter discusses the necessity of examining your system's security and reliability requirements as early in the software design process as feasible. Furthermore, the research dives into the tradeoffs between automation and human involvement, examining the advantages and drawbacks of each technique. It covers the tradeoffs in allocating resources between creating new features and establishing strong systems, finding a balance between short-term aims and long-term sustainability.

The study's results may help Site Reliability Engineers, software architects, and developers make educated judgements when faced with design compromises. By understanding the possible tradeoffs and their ramifications, practitioners may more successfully negotiate the intricacies of SRE, resulting in more robust and dependable systems. We begin by discussing the relationship between system restrictions and product features, followed by two examples a payment processing service and a microservices framework that show some typical security and reliability

choices. We finish by discussing the natural temptation to postpone security and reliability work, as well as how early investment in security and reliability may lead to sustained project pace.

In site reliability engineering (SRE), design tradeoffs relate to the sacrifices and decisions that must be made while developing and executing systems to maintain their dependability. SRE is a field that focuses on maintaining and enhancing the reliability of large-scale distributed systems, and it often necessitates challenging choices to balance competing needs and restrictions.

Design compromises in SRE originate from the requirement to achieve high reliability while also taking into account considerations such as performance, scalability, cost, and maintainability. Optimizing a system for optimum dependability, for example, may include more redundancy and fault-tolerance measures, which might raise complexity and expense. Prioritizing cost-efficiency, on the other hand, may imply abandoning certain reliability aspects or choosing lower-cost components that may be less dependable [1]–[3].

Additionally, design tradeoffs in SRE may include taking into account the influence on other areas of the system or organisation. Implementing tougher security measures, for example, to improve dependability may involve more overhead and may degrade system speed. It is critical to balance these tradeoffs in order to maintain optimum dependability while satisfying other needs and limits. Finally, SRE design compromises need thorough study and analysis of many elements in order to find a balance between dependability and conflicting goals. SRE teams can build and run systems that are robust, scalable, cost-effective, and maintainable by properly understanding and managing these tradeoffs, allowing organisations to deliver a dependable and smooth user experience.

## **DISCUSSION**

So, you intend to create a (software) product! In this challenging trip, you'll have a lot to think about, from developing high-level strategies to delivering code. Typically, you'll begin with a broad notion of what the product or service will perform. This may be a high-level idea for a game or a collection of high-level business needs for a cloud-based productivity solution. You'll also create high-level ideas on how to finance the service providing. Additional needs and restrictions on the design and execution of the application tend to arise as you go through the design phase and your thoughts about the form of the product become more detailed. There will be particular needs for the product's functioning as well as general limits such as development and operating expenses. You'll also have security and reliability needs and constraints: your service will most likely have availability and reliability requirements, and you may have security requirements for securing sensitive customer data handled by your application. Some of these criteria and limitations may be in conflict, and you'll need to make choices to strike the proper balance.

### **Design Goals and Requirements**

Your product's feature needs will have dramatically different characteristics than your security and reliability requirements. Let's take a deeper look at the different sorts of requirements you'll encounter while creating a product.

### **Requirements for Features**



Feature requirements, also known as functional requirements, highlight a service's or application's core purpose and specify how a user may complete a certain job or meet a specific demand. They are often represented in terms of use cases, user stories, or user journeys, which are sequences of encounters between a user and a service or application. The subset of feature needs that are vital to the product or service is known as critical requirements. You don't have a viable product if a design does not meet a key requirement or important user narrative. Typically, feature needs are the key motivators for your design selections. After all, you're attempting to create a system or service that meets a certain set of requirements for the people you have in mind. You often have to choose between competing needs. With this in mind, it is helpful to differentiate between critical and non-critical feature needs.

A handful of criteria usually apply to the whole application or service. These criteria are often absent from user stories or particular product needs. Instead, they are mentioned once in centralized requirements documents, or are even assumed implicitly. Here's an example:

1. All views/pages of the application's web UI must.
2. Follow common visual design guidelines.
3. Adhere to accessibility guidelines.
4. Have a footer with links to privacy policy and ToS (Terms of Service).

### Nonfunctional Requirements

Several requirement types concentrate on generic system traits or behaviours rather than particular behaviours. These nonfunctional needs are pertinent to our primary concern security and dependability. As an example:

- a. What are the only conditions under which someone (an external user, customer service representative, or operations engineer) may have access to certain data?
- b. What are the service level objectives (SLOs) for metrics like uptime or response latency in the 95th and 99th percentiles? How does the system react when the load exceeds a specific threshold?[4]–[6]

When balancing needs, it might be useful to examine requirements in areas outside than the system itself, since decisions in other areas can have a major influence on core system requirements.

Among these wider areas are the following:

- 1. Development efficiency and velocity:** How efficiently can developers iterate on new features given the implementation language, application frameworks, testing procedures, and build processes? How well can developers comprehend, alter, and troubleshoot existing code?
- 2. Deployment velocity:** How long does it take for a feature to be created before it is made accessible to users/customers?

### Balancing Requirements

Because a system's characteristics that address security and reliability issues are essentially emergent, they often interact with one other and with how feature needs are implemented. As a

consequence, it may be difficult to think critically about security and dependability considerations on their own. Because security and reliability are emergent concepts, design decisions relating to them are frequently quite fundamental. They are comparable to fundamental architectural decisions like whether to use a relational or NoSQL database for storage or a monolithic or microservices architecture. An existing system that wasn't created with security and dependability in mind from the beginning often has trouble being "bolted on" with these features. A system is more likely to have poorer availability and be more vulnerable to faults with security implications if the interfaces between its components are poorly defined and difficult to comprehend. Testing and strategic bug-fixing cannot alter it.

It is sometimes necessary to make considerable design modifications, extensive refactorings, or even partial rewrites in order to accommodate security and reliability requirements in an existing system. These changes may be highly costly and time-consuming. Additionally, such modifications could need to be performed quickly in response to a security or reliability event. However, making large design modifications to a deployed system quickly carries a high risk of introducing new faults. Therefore, from the first planning stages of a software project, it is crucial to take security and reliability needs into account and the accompanying design compromises. If your company has security and SRE teams, they should be included in these talks.

An illustration of the kind of compromises you could have to make is provided in this section. This example includes certain sections that go rather in-depth into technical aspects, which may or may not be significant in and of itself. For this example, it is also not necessary to discuss any of the compliance, regulatory, legal, or commercial factors that are taken into account when creating and operating payment processing systems. The aim is to highlight the intricate interdependencies between needs. In other words, the design process for a system with intricate security and dependability needs is the emphasis, not the specifics of securing credit card information.

### **Using a third-party service provider to handle sensitive data**

The best method to address security issues with sensitive data is often to avoid holding it in the first place (for more information on this subject, see Chapter 5). You may be able to make sure that sensitive data never enters your systems, or at the very least, build the systems such that the data isn't stored there permanently.<sup>4</sup> You may pick from a variety of commercial payment service APIs to connect with the application, and you can leave the vendor in charge of processing payment information, payment transactions, and associated issues (such fraud prevention measures).

### **Benefits**

- 1) Using a payment service may, depending on the circumstances, minimise risk and the extent to which you need to develop internal knowledge to handle concerns in this area, relying instead on the provider's experience:
- 2) Sensitive data is no longer stored on your systems, which lowers the possibility that a flaw in your procedures or systems may cause a data breach. Of course, the data of your users may still be compromised if the third-party provider were compromised.

- 3) Your contractual and compliance duties under payment industry security standards may be streamlined depending on the particulars and relevant criteria.
- 4) To safeguard the data stored at rest in your system's data storage, you don't need to construct and maintain infrastructure. This might reduce current operational and development work substantially.
- 5) A lot of third-party payment service providers provide services for payment risk assessment and countermeasures against fraudulent transactions. These features may allow you to lower your risk of payment fraud without having to develop and manage the underlying infrastructure yourself[7]–[9].

On the other hand, using a third-party service provider comes with its own expenses and dangers.

### **Costs and nontechnical risks**

1. The provider will undoubtedly impose surcharges. Your decision will probably be influenced by the amount of transactions; over a certain level, processing transactions internally is usually more cost-effective.
2. The technical cost of depending on a third party dependence should also be taken into account. Your team will need to learn how to use the vendor's API, and you may need to keep track of changes/releases of the API according to the vendor's timetable.

### **Reliability risks**

You introduce a new reliance to your application in this example, a third-party service by outsourcing payment processing. There are often more failure possibilities introduced by extra dependencies. These failure types could be partly beyond your control if there are third-party dependencies. For instance, if the payment provider's service is unavailable or not accessible over the network, your user narrative "user can buy their chosen widgets" may not operate. The importance of this risk relies on how well the payment provider follows the SLAs you have with that supplier.

By providing redundancy to the system (see Chapter 8), in this example by include a backup payment provider to which your service may fail over, you can mitigate this risk. The two payment providers most likely have distinct APIs, so you must build your system to be able to communicate with them. This adds to the technical and operational expenses and increases the risk of defects or security breaches. Additionally, you might reduce the dependability risk by implementing contingency plans on your end. If the payment service is unavailable, for instance, you may add a queuing mechanism to the communication channel with the payment provider to buffer transaction data. By doing this, the "purchase flow" user narrative would be able to continue even if the payment service went down.

The message queuing technique, however, adds complexity and could generate new failure possibilities. You run the danger of losing transactions if the message queue is not built to be dependable (it only stores data in volatile memory, for instance). Generally speaking, reliability problems and hidden flaws may exist in subsystems that are only sometimes and under unusual conditions used. A more dependable message queue solution is an option. This most likely entails permanent disc storage or an in-memory storage system dispersed over many physical locations, which adds complexity once again. Even in rare circumstances, saving the data on disc brings

back the issues with storing sensitive data (risk of breach, compliance issues, etc.) that you were initially attempting to avoid. It is particularly challenging to implement a retry queue that uses persistent storage in this situation since certain payment data isn't even permitted to reach the disc. In light of this, you may need to take into account attacks (especially insider attacks) that deliberately damage the connection with the payment provider in order to start local queueing of transaction data, which might later be compromised[10].

### **Security risks**

The decision to depend on a third-party service during design also brings up urgent security issues. First off, you're giving a third-party vendor access to confidential client information. You should choose a vendor whose security attitude is at least comparable to your own, and you should carefully consider vendors both while choosing them and moving forward. This is not a simple process, and there are many contractual, governmental, and liability factors that go beyond the purview of this book and that you should discuss with your legal counsel.

Second, linking a vendor-supplied library into your application could be necessary in order to integrate with the vendor's service. This raises the possibility that a flaw in that library or one of its transitive dependencies might lead to a flaw in your systems. By sandboxing the library<sup>5</sup> and being ready to swiftly release new versions of it, you may think about reducing this risk. By working with a vendor that doesn't force you to integrate a proprietary library into your service, you may substantially allay this worry. If the company provides its API using an open protocol like REST+JSON, XML, SOAP, or gRPC, proprietary libraries may be avoided.

In order to interface with the vendor, your web application client may need to incorporate a JavaScript library. By doing this, you may prevent the temporary passage of payment information via your systems and enable the direct transmission of payment information from a user's browser to the web service of the supplier. The vendor's library code executes with full rights in the web origin of your application, which creates comparable issues as adding a server-side library.<sup>6</sup> Your application might be hacked by a flaw in that code or a server compromise that affects the server hosting that library. You may think about sandboxing payment-related functionality in a different web origin or sandboxed iframe to reduce that risk. However, this strategy necessitates the use of a secure cross-origin communications method, which adds complexity and more potential points of failure. The payment provider may also provide an HTTP redirect-based integration, although this may lead to a less seamless user experience.

In areas of domain-specific technical expertise, design decisions relating to nonfunctional requirements can have fairly broad repercussions. For example, while debating a tradeoff involving reducing the risks associated with handling payment data, we ended up thinking about issues that are deeply rooted in web platform security. We also ran across contractual and regulatory issues along the route.

### **Reliability and Security Benefits of Software Development Frameworks**

Developers use the framework because it streamlines application development and automates typical tasks, making their everyday job simpler and more productive. A solid and widely used framework with built-in dependability and security features is a win-win situation. The framework offers a shared feature surface on which security engineers and SREs may develop

new functionality, opening up chances for further automation and quickening the pace of the project as a whole.

However, since this architecture automatically addresses typical security and reliability problems, building on it leads to systems that are intrinsically more secure and dependable. Additionally, it greatly improves the efficiency of security and production readiness reviews. For example, if a software project's continuous builds and tests are green, meaning that its code passes conformance checks at the framework level, you can be sure that it is not impacted by the common security issues that the framework has already addressed. The framework's deployment automation similarly makes sure that the service meets its SLA by pausing releases when the error budget is depleted; for more information, see Chapter 16 of the SRE workbook. SREs and security engineers may utilise their time to concentrate on more intriguing design-level issues. Finally, anytime an application is rebuilt (with the most recent dependencies) and deployed, bug fixes and enhancements in the framework's code are immediately transmitted to the application.

## CONCLUSION

Since security and reliability are largely emergent features of the whole development and operations process, designing and developing safe and reliable systems is not a simple task. This task requires considering a wide range of quite complicated issues, many of which first seem to have nothing to do with meeting the main feature needs of your service. Numerous security, reliability, and feature trade-offs will need to be made throughout the design phase. These compromises often seem to be at odds with one another at first. Early on in a project, it may seem appealing to put off dealing with these problems in favor of "dealing with them later"; yet, doing so often entails considerable costs and project risks since dependability and security are essential components once your service is operational. If your service is unavailable, you risk losing customers; if your service is compromised, you'll need all hands-on deck to react. But it is often feasible to meet all three of these qualities with thoughtful planning and design. Additionally, you may achieve this for a very little up-front extra cost and often with less overall engineering work during the system's lifespan.

## REFERENCES

- [1] J. Derboven, R. Voorend, and K. Slegers, "Design trade-offs in self-management technology: the HeartMan case," *Behav. Inf. Technol.*, 2020, doi: 10.1080/0144929X.2019.1634152.
- [2] D. Fogli, A. Piccinno, S. Carmien, and G. Fischer, "Exploring design trade-offs for achieving social inclusion in multi-tiered design problems," *Behav. Inf. Technol.*, 2020, doi: 10.1080/0144929X.2019.1634153.
- [3] A. Schulz, H. Wang, E. Grinspun, J. Solomon, and W. Matusik, "Interactive exploration of design trade-offs," *ACM Trans. Graph.*, 2018, doi: 10.1145/3197517.3201385.
- [4] A. Andrews, E. Mancebo, P. Runeson, and R. France, "A framework for design tradeoffs," *Softw. Qual. J.*, 2005, doi: 10.1007/s11219-005-4252-z.
- [5] W. L. Luyben, "Economic trade-offs in acrylic acid reactor design," *Comput. Chem. Eng.*, 2016, doi: 10.1016/j.compchemeng.2016.06.005.
- [6] L. L. Bailey, J. E. Hines, J. D. Nichols, and D. I. MacKenzie, "Sampling design trade-offs

- in occupancy studies with imperfect detection: Examples and software,” *Ecol. Appl.*, 2007, doi: 10.1890/1051-0761(2007)017[0281:SDTIOS]2.0.CO;2.
- [7] M. A. Rau, W. Keesler, Y. Zhang, and S. Wu, “Design Tradeoffs of Interactive Visualization Tools for Educational Technologies,” *IEEE Trans. Learn. Technol.*, 2020, doi: 10.1109/TLT.2019.2902546.
- [8] P. P. Pande, C. Grecu, M. Jones, A. Ivanov, and R. Saleh, “Performance evaluation and design trade-offs for network-on-chip interconnect architectures,” *IEEE Trans. Comput.*, 2005, doi: 10.1109/TC.2005.134.
- [9] G. Fischer, “Identifying and exploring design trade-offs in human-centered design,” in *Proceedings of the Workshop on Advanced Visual Interfaces AVI*, 2018. doi: 10.1145/3206505.3206514.
- [10] C. Xu, S. Sugiura, S. X. Ng, P. Zhang, L. Wang, and L. Hanzo, “Two Decades of MIMO Design Tradeoffs and Reduced-Complexity MIMO Detection in Near-Capacity Systems,” *IEEE Access*. 2017. doi: 10.1109/ACCESS.2017.2707182.

## DESIGN FOR LEAST PRIVILEGE

Dr. Shrishail Anadinni\*

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:shrishail@presidencyuniversity.in

---

### ABSTRACT:

*Whether access is provided by systems or by people, the concept of least privilege states that users should only have the access necessary to complete a job. The best time to implement these limitations is at the outset of the development lifecycle, during the design stage of brand-new functionality. Unnecessary privilege increases the chance of errors, faults, or breach and increases the cost of mitigating or eliminating security and reliability issues in an operational system. In this chapter, we go through risk-based access classification and recommended practices for enforcing least privilege. An illustration of configuration distribution shows the trade-offs that these techniques involve in actual contexts. After describing a framework for authentication and authorization policies, we go in-depth on sophisticated authorization controls that may reduce the danger of hacked accounts as well as the risk of errors made by on-call engineers. In order to draw a conclusion, we acknowledge the difficulties and compromises that may occur while designing with least privilege. In a perfect world, system users are well-intentioned and carry out their job duties flawlessly securely and without making any mistakes. Sadly, the truth is rather different. Engineers might commit fraud, have their accounts hacked, or create errors on purpose. For these reasons, it's critical to design systems with the least amount of privilege possible. Systems should only provide users access to the information and services they need to do their tasks.*

**KEYWORDS:** *Breakglass, Isolation, Minimum Privilege, Segmentation.*

---

### INTRODUCTION

In site reliability engineering (SRE), designing for least privilege entails putting security mechanisms in place that guarantee people, systems, and processes are only given the minimal access rights required to carry out their assigned duties. A basic security idea called the "principle of least privilege" seeks to reduce the possible harm that compromised or malevolent entities may be able to do[1]–[3]. Designing for least privilege in the context of SRE involves:

#### **RBAC (role-based access control)**

Organisations may design and enforce granular access rights based on roles and responsibilities by using RBAC. The attack surface may be reduced and the possible effects of security breaches can be mitigated by restricting access to critical resources and only giving rights to the precise roles that need them.

### **Isolation and segmentation**

By limiting the breadth of the breach, a segmented architecture that isolates various parts and services may lessen the effect of a security event. The danger of lateral movement and unauthorized access may be reduced by compartmentalizing resources and restricting connectivity between them.

### **The least privilege principle for processes:**

Applying the least privilege concept guarantees that processes operating on the system are only given the rights required to complete their duties. This strategy guards against unauthorized or compromised processes from potentially abusing their privileged capabilities.

### **Regular access reviews**

Regular access audits and reviews assist find and eliminate privileges that have been given to people or processes that are too generous or unneeded. This procedure lowers the possibility of unauthorized access while ensuring that privileges are consistently in line with job requirements.

### **Mechanisms for privilege escalation**

Implementing stringent controls and privilege escalation procedures may assist restrict the duration and extent of elevated privileges when privileged access is necessary. This lessens the possible effects of privilege breach and helps prevent unauthorized access. A thorough grasp of the system needs, roles, and responsibilities, as well as possible security risks and vulnerabilities, is necessary when designing for least privilege in SRE. Organisations may lower the attack surface, lessen the effect of security events, and improve the overall security and dependability of their systems by adhering to this concept and putting the right security measures in place.

## **DISCUSSION**

Companies often want to assume the best about their engineers and depend on them to do impossible jobs perfectly. This is not a realistic goal. Think about the harm you might do to your company if you decided to do anything bad as a mental exercise. How could you respond? What method would you use? Would you be found out? Could you obfuscate your actions? What's the worst mistake you or someone with equal access could make, even if your intentions were good? How many ad hoc, manual commands performed by you or your colleagues might result in an outage or make it worse while troubleshooting, reacting to an outage, or carrying out emergency response?

We must presume that any potential negative action or event is feasible since we cannot depend on human perfection. Therefore, we advise constructing the system to reduce or do away with the effects of these undesirable deeds. Even though we typically have faith in the people using our systems, we still need to set certain restrictions on their access rights and level of credential trust. Problems can and will arise. People will commit errors, misuse commands, become vulnerable, and fall for phishing emails. It is absurd to anticipate perfection. In other words, hoping is not a strategy, to paraphrase an SRE principle.

### **Terminology and Concepts**

Let's establish workable definitions for a few key terms of art used in the sector and at Google before we get into best practices for creating and running an access control system.



### **Minimum Privilege**

In the security business, the idea of least privilege is a wide one that is well-established. The system that gives the least privilege required for every given job or action route may be built using the high-level best practices in this chapter. This objective pertains to the people, computers, and automated tasks that make up a distributed system. All system levels for authentication and authorization should adhere to the principle of least privilege. Our suggested method strives to prevent users from having ambient power for instance, the ability to log in as root as far as is physically practicable. In particular, it resists giving implicit access to tools as seen in Worked Example: Configuration Distribution.

### **Zero Trust Networking**

The idea that a user's network location inside the company's network doesn't offer any privileged access is the foundation of the design concepts we describe. This is known as zero trust networking. For instance, connecting from another location on the internet is just as effective as hooking into a network connection in a conference room. Instead, access is granted by a system based on a mix of user and device credentials that is, on what we know about the user and the device. Through its BeyondCorp programme, Google has successfully established a zero trust networking paradigm on a broad scale[4]–[6].

### **Zero Touch**

In order to transition to what we refer to as Zero Touch interfaces, the SRE organisation at Google is trying to expand on the idea of least privilege via automation. By eliminating direct human access to production roles, these interfaces such as Zero Touch Production (ZTP), which is discussed in Chapter 3, and Zero Touch Networking (ZTN) are intended to make Google safer and decrease downtime. Instead, tools and automation that make predictable and regulated modifications to the production infrastructure provide people indirect access to the production process. Significant automation, new secure APIs, and robust multi-party approval mechanisms are all necessary for this strategy.

### **Classifying Access Based on Risk**

Tradeoffs are a part of any risk reduction approach. Increasing engineering effort, process adjustments, operational work, or opportunity cost may be necessary to reduce the risk brought on by human actors. This might result in productivity tradeoffs. By clearly defining your objectives and setting priorities, you may reduce these expenditures. The composition of your access may vary significantly depending on the nature of your system since not all data or activities are created equal. As a result, you shouldn't defend each access point equally. You must categorise access based on impact, security risk, and/or criticality in order to apply the best restrictions and prevent an all-or-nothing mindset. For instance, you'll probably need to manage access to various data kinds (publicly accessible data, corporate data, user data, and cryptographic secrets) in different ways. Similar to how you should handle service-specific read APIs differently from administrative APIs that provide data deletion capabilities.

In order for individuals to create systems and services that "speak" your categories, they must be precisely defined, consistently applied, and widely understood. Depending on the size and complexity of your system, your classification framework will change. You could simply require two or three categories that depend on ad hoc labelling, or you might need a sophisticated and

automated system for categorizing different system components (such as API groups and data types) in a central inventory. These categories could apply to data repositories, APIs, services, or other entities that users might access while working. Make sure your framework can manage the most crucial components of your systems.

Consider the controls in place for each categorization after you've built a foundation for it. You must take into account the following factors:

1. Who is to be allowed access?
2. How strictly should access be limited?
3. The user needs read-only or read/write access.
4. What safeguards are in place for the infrastructure?

For instance, a business could need the three categories public, sensitive, and very sensitive, as indicated in #example\_access\_classifications\_based\_on. Depending on the amount of harm that improperly allowed access may cause, that organisation could classify security measures as low, medium, or high risk.

	Description	Read access	Write access	Infrastructure access <sup>1</sup>
<b>Public</b>	Open to anyone in the company	Low risk	Low risk	High risk
<b>Sensitive</b>	Limited to groups with business purpose	Medium/high risk	Medium risk	High risk
<b>Highly sensitive</b>	No permanent access	High risk	High risk	High risk

**Figure 1: Represents Example access classifications based on risk.**

Building an access architecture that will allow you to implement suitable controls with the ideal ratio of productivity, security, and dependability should be your aim. All data access and operations should follow the principle of least privilege. Let's talk about how to design your systems using the ideas and controls for least privilege, building on this fundamental framework.

### BEST PRACTICES

We suggest a number of recommended practices, which are listed below, when adopting a least privilege model.

#### Small functional APIs

This comment illustrates how tiny, flexible tools are important to Unix culture. The authors' advice still holds true more than 40 years later since current distributed computing developed from the single time-shared computer systems of the 1970s into planet-wide network-connected distributed systems. One may modify this quotation to read, "Make each API endpoint do one thing well," to better fit the contemporary computing landscape. Avoid open-ended interactive interfaces when designing systems with security and dependability in mind; instead, focus on

creating modest, useful APIs. With this method, you may use the time-tested security concept of least privilege and only provide users the rights required to carry out a certain task[7]–[9].

What do we precisely mean by API? Every system has an API, which is just the graphical user interface it offers. The POSIX API<sup>3</sup> and the Windows API<sup>4</sup> are examples of very big APIs. Memcached and NATS are examples of very small APIs. The World Clock API, TinyURL, and the Google Fonts API<sup>8</sup> are examples of extremely small APIs. The total number of methods to query or change the internal state of a distributed system is what is meant when we speak about the API of that system. This chapter focuses on how to develop and maintain secure systems by providing API endpoints with a few well-defined primitives. API design has been extensively studied in computer literature. Instead of an API that takes a programming language, the input you assess, for instance, may be CRUD (Create, Read, Update, and Delete) activities on a certain ID.

Pay close attention to the administrative API in addition to the user-facing API. For the dependability and security of your application, the administrative API is just as crucial—if not more so. When accessing these APIs, typos and errors might disclose vast quantities of data or cause catastrophic disruptions. As a consequence, hostile actors find administrative APIs to be some of the most alluring attack surfaces. Compared to user-facing APIs, administrative APIs might be quicker and simpler to alter since they are solely used by internal users and tools. However, there will still be a cost to modifying any API once your internal users and tools start developing on it, therefore we advise carefully examining its design.

Administrative APIs include the following:

- i. APIs for setup and teardown, such those needed to create, install, and update software or provide the container or containers it runs in
- ii. APIs for routine maintenance and emergencies, such as administrator access to erase user data or state that has been damaged or to restart problematic processes

### **Does an API's size affect access and security in any way?**

Take a look at a well-known illustration: the POSIX API, one of our earlier illustrations of a very broad API. This API is well-liked since it is adaptable and well known. Its most frequent use as a production machine administration API includes restarting daemons, altering configuration files, and installing software packages. Users often configure and maintain classic Unix<sup>10</sup> hosts using an interactive OpenSSH connection or tools that script against the POSIX API. Both methods make the whole POSIX API available to the caller. Controlling and monitoring a user's behavior during that interactive session might be challenging. This is particularly true if the user is trying to bypass the controls on purpose or if the connected workstation has been hacked. The POSIX API<sup>11</sup> allows you to restrict the rights supplied to the user, but this requirement is a fundamental flaw of providing a very wide API.

Instead, it would be preferable to divide up this enormous administrative API into more manageable chunks. Then, you can use the concept of least privilege to only provide authorization for the exact action(s) that each caller specifically requests. It is important to distinguish the OpenSSH API from the disclosed POSIX API. By utilising `git-shell`, for instance, it is feasible to use OpenSSH and its authentication, authorization, and auditing (AAA) controls without disclosing the complete POSIX API[10].

## Breakglass

A breakglass mechanism, so named from fire alarm pulls that advise the user to "break glass in case of emergency," grants access to your system in an emergency and fully avoids your authorization system. For bouncing back from unanticipated events, this may be helpful. See Graceful Failure, Breakglass Mechanisms, and Diagnosing Access Denials for further information.

## CONCLUSION

In conclusion, Site Reliability Engineering (SRE) is essential for ensuring the security and integrity of systems. This principle is known as designing for least privilege. The least privilege principle states that users or processes should only be given the minimal amount of power required to carry out their intended duties. SRE teams may greatly decrease the attack surface and the negative effects of security breaches by adhering to this approach. Designing for least privilege entails carefully assessing and restricting the access and permissions given to various system components and users. It requires a full comprehension of the needs and dependencies of the system, as well as an evaluation of possible risks and weaknesses. Limiting administrative access, imposing stringent authentication and authorization procedures, creating strong access controls, and routinely evaluating and updating permissions are all examples of how least privilege might be implemented in SRE. Strong monitoring and auditing capabilities should be used in conjunction with the least privilege principal implementation in order to spot possible violations and unauthorized access attempts and take appropriate action. SRE teams may improve the system's overall security posture, lessen the risk of unauthorized access, and lessen the possible impact of security events by designing for least privilege. This strategy not only assists in safeguarding sensitive information and assets but also promotes a culture of security and accountability inside the company.

## REFERENCES

- [1] K. Buyens, R. Scandariato, and W. Joosen, "Least privilege analysis in software architectures," *Softw. Syst. Model.*, 2013, doi: 10.1007/s10270-011-0218-8.
- [2] U. D. Upadhyay, A. F. Cartwright, V. Goyal, E. Belusa, and S. C. M. Roberts, "Admitting privileges and hospital-based care after presenting for abortion: A retrospective case series," *Health Serv. Res.*, 2019, doi: 10.1111/1475-6773.13080.
- [3] N. Roessler *et al.*, "µSCOPE: A methodology for analyzing least-privilege compartmentalization in large software artifacts," in *ACM International Conference Proceeding Series*, 2021. doi: 10.1145/3471621.3471839.
- [4] A. Barth, A. P. Felt, P. Saxena, and A. Boodman, "Protecting Browsers from Extension Vulnerabilities," *Ndss*, 2010.
- [5] Aramark, *Corp. Philanthr. Rep.*, 2019, doi: 10.1002/cprt.30318.
- [6] Z. Wang, H. Pang, and Z. Li, "Access control for Hadoop-based cloud computing," *Qinghua Daxue Xuebao/Journal Tsinghua Univ.*, 2014.
- [7] G. Ferreira, L. Jia, J. Sunshine, and C. Kastner, "Containing malicious package updates in npm with a lightweight permission system," in *Proceedings - International Conference on*

- Software Engineering*, 2021. doi: 10.1109/ICSE43902.2021.00121.
- [8] S. Chen, “###Chrome Extensions : Threat Analysis and Countermeasures,chrome,” *Ndss*, 2012.
- [9] S. Steiner, D. C. De Leon, and A. A. Jillepalli, “Hardening web applications using a least privilege DBMS access model,” in *ACM International Conference Proceeding Series*, 2018. doi: 10.1145/3212687.3212863.
- [10] S. Kaya, “From needs analysis to development of a vocational English language curriculum: A practical guide for practitioners,” *J. Pedagog. Res.*, 2021, doi: 10.33902/jpr.2021167471.

**AUTHORIZATION DECISIONS, TRADEOFFS AND TENSIONS****Ms. Hireballa Sangeetha\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA,  
Email Id:sangeethahm@presidencyuniversity.in

**ABSTRACT:**

*Designing safe and dependable systems involves making judgements about authorization as well as dealing with trade-offs and tensions. Determining the permissions and access privileges given to people or processes inside a system is part of the authorization process. Making permission choices, however, requires taking into account a number of conflicts and tradeoffs. The trade-off between security and usability in authorization is one of the main issues. Although rigorous access restrictions and constrained permissions increase security, they may also reduce user productivity and add administrative burden. User demands, system requirements, and possible dangers must all be carefully considered in order to strike the correct balance between security and usability. Additionally, conflicts between security and performance often need to be taken into account when making permission choices. Complex permission checks may increase costs and have an adverse effect on system performance. It's critical to strike the ideal balance between preserving system performance and establishing safe access management.*

**KEYWORDS:** Authentication, Complexity, Decentralization, Tradeoffs.

**INTRODUCTION**

In software engineering, authorization choices entail deciding what actions and resources users or processes are permitted to access. These choices are essential for upholding effective governance, safeguarding sensitive data, and ensuring security. However, permission choices also bring up difficulties and tradeoffs that must be properly taken into account. The tension between complexity and maintainability in authorization is another trade-off. Fine-grained access control implementation may improve security, but it can also provide complicated authorization rules that are challenging to administer and maintain. Although it may increase maintainability, security may be compromised by simplifying permission processes. It's crucial to determine the ideal amount of complexity that guarantees both security and administration simplicity.

In general, authorization choices, trade-offs, and tensions are important factors to take into account while creating safe and trustworthy systems. Organisations may develop efficient authorization methods that strike a balance between security, usability, maintainability, and performance by carefully assessing the tradeoffs and tensions. Conflicts between centralization and decentralization can exist in authorization choices. In addition to providing consistency and control, centralized authorization solutions may also lead to single points of failure and scalability issues. While offering flexibility and scalability, decentralized authorization techniques might be more challenging to administer and coordinate. Depending on the needs of the system, the organizational structure, and the risk tolerance, the right amount of centralization

or decentralization must be chosen. The balance between usability and security is one trade-off in authorization choices.

By restricting access to critical resources, stricter authorization procedures and controls may improve security, but they may also add complexity and discomfort for users. In order to avoid having a detrimental effect on productivity and user happiness, it's essential to strike the correct balance between robust security measures and a user-friendly experience. There is also another conflict between flexibility and control. Giving users greater control over their authorization privileges might help them do their jobs more quickly, but it also raises the possibility of abuse or unauthorized access. The implementation of severe rules and limitations, however, might impede productivity and creativity while offering more control. Understanding the unique requirements of the system and its users and putting the proper controls in place as a result are necessary for striking the correct balance [1]–[3].

Conflicts between centralization and decentralization can exist in authorization choices. It may be simpler to enforce regulations and keep track of access with the help of centralized authorization management, which can provide a single and consistent approach. However, it could also add administrative burden and single points of failure. Decentralized authorization may empower specific teams or departments and disperse decision-making, but it may also lead to inconsistencies and make it more difficult to manage access throughout the organisation. The system's context and particular needs determine the extent of centralization or decentralization that is suitable.

Additionally, there is a conflict in permission choices between granularity and simplicity. Fine-grained permissions are made possible by granular access restrictions, ensuring that users only have access to the resources they need. However, it might be complicated and challenging to handle a large number of granular permissions. Broader permissions may make authorization choices easier to handle, but they can raise the danger of unauthorized access or power escalation. It is important to properly balance tradeoffs and tensions while making authorization choices in software engineering. Designing an effective and efficient authorization system requires striking a balance between security and usability, flexibility and control, centralization and decentralization, and granularity and simplicity. While still addressing the demands of the organisation and its users, understanding the precise requirements and limitations of the system and regularly assessing and adjusting the authorization method may help avoid risks and provide adequate access restrictions.

## **DISCUSSION**

### **A Policy Framework for Authentication and Authorization Decisions**

An organized method for creating and implementing access restrictions inside a system is provided by a policy framework for authentication and authorization decisions. It creates a collection of standards, regulations, and processes that control the authentication and authorization procedure, making sure that only entities that have been authenticated and granted access to resources may do so. The essential elements of a policy framework for choices about authentication and authorization are as follows:

#### **Policy Definition**

The authentication and authorization system's aims, objectives, and requirements must be clearly stated. The categories of people, resources, and behaviors that need to be safeguarded must be identified, and any legal or compliance requirements must be taken into account.

### **Authentication Policies**

Decide on standards for verifying people or processes trying to access the system. This can include requiring strong passwords, using several login methods, or integrating with other identity services.

### **Authorization Policies**

Establish policies and techniques for access control to decide which resources and activities people or processes may access. The use of attribute-based access control (ABAC), role-based access control (RBAC), and other access control models may be involved in this.

### **Principle of Least Privilege**

Use the least privilege concept to ensure that users or processes only get the permissions required to complete their responsibilities. The potential harm caused by unauthorized access or privilege escalation is lessened as a result [4]–[6].

### **Policy Enforcement**

Implement controls to ensure that the stated rules for authentication and authorization are followed. Implementing firewalls, access control lists (ACLs), or other security tools that verify and uphold access choices may be necessary for this.

### **Auditing and Monitoring**

Create systems for keeping an eye on and auditing authentication and authorization processes. This entails tracking access attempts, spotting abnormalities or suspect activity, and carrying out regular audits to guarantee policy observance.

### **Policy Review and Update**

Review and update the authentication and authorization rules often to reflect changing organizational demands, technological developments, and security requirements. This involves assessing fresh access control techniques or innovations that might boost safety and effectiveness.

### **Training and Awareness**

To inform users, administrators, and developers about authentication and authorization rules and best practices, provide training and awareness programs. This ensures that the policies are followed and implemented correctly.

### **Compliance and Governance**

Align the rules for authentication and authorization with any applicable legal mandates and organizational governance structures. This involves making sure that company rules, industry standards, and data protection laws are followed. Organisations may develop reliable access controls that safeguard sensitive resources, reduce security risks, and guarantee regulatory



compliance by putting in place a thorough policy framework for authentication and authorization choices.

### **Using Advanced Authorization Controls**

An established method of putting an authorization decision into practice is an access control list for a specific resource. The most basic ACL is a string that matches the authenticated role, often associated with some kind of grouping—for instance, a set of roles, like "administrator," that grows to a longer list of roles, like usernames. The authenticated role must be a member of the ACL in order for the service to analyse the incoming request. Multi-factor authentication (MFA) and multi-party authorization (MPA), for example, need more complicated authorization codes (see Advanced Controls for more information on MFA and MPA). Additionally, while developing authorization procedures, certain organisations may need to take into account their own regulatory or contractual obligations.

Correctly implementing this code may be challenging, and if several services each implement their own authorization logic, the complexity of the code will grow quickly. In our experience, frameworks like the AWS or GCP Identity & Access Management (IAM) services are useful for separating the complexity of authorization choices from fundamental API architecture and business logic. For internal services, Google also makes considerable use of a variant of the GCP permission system. Our code can do basic checks (such "Can X access resource Y?") and compare those tests to externally specified policies thanks to the security policy framework. We just need to make a little modification to the necessary policy configuration file if we need to add extra permission controls to a certain activity. This tiny overhead has a huge impact on functionality and speed.

### **Identifying and Avoiding Potential Pitfalls**

It's challenging to create a sophisticated authorization policy language. The aim won't be achieved if the policy language is too straightforward, and you'll wind up having authorization choices scattered across the framework's policy and the main codebase. It may be exceedingly challenging to rationally justify a policy if the terminology is excessively broad. You may use best practices for software API design, particularly an iterative design approach, to allay these worries, but we advise continuing cautiously to stay away from both of these extremes. Pay close attention to the authorization policy's shipping to (or inclusion in) the binary. Independent of the binary, you could wish to alter the authorization policy, which is likely to become one of the configuration elements with the highest security risk.

The policy choices that will be encoded in this language will need cooperation from application developers. To strike the right balance between security and functionality, it will typically take cooperation between application developers implementing the administrative APIs and security engineers and SREs with domain-specific knowledge about your production environment, even if you avoid the pitfalls described here and develop an expressive and understandable policy language.

### **Temporary Access**

By allowing momentary access to resources, you may reduce the danger of an authorization decision. This approach is often helpful when you wish to offer the fewest privileges feasible with the available technology but fine-grained controls are not accessible for every activity[7]–

[9]. Temporary access may be granted in a planned and controlled manner (such as during on-call rotations or via expiring group memberships) or on-demand, when users voluntarily seek access. A request for multi-party authorization, a business rationale, or another authorization control may all be combined with temporary access. Temporary access also provides a reasonable starting point for auditing since you can clearly see who has access at any given moment thanks to transparent logging. In order to prioritize and gradually lower these requests, it also gives information about the locations where temporary access happens. Additionally, temporary access lessens ambient authority. This is one reason why administrators prefer sudo or "Run as Administrator" over using the Unix user root or Windows Administrator accounts: the less rights you have, the less likely you are to unintentionally submit a command that deletes all the data!

### Proxies

You may fall back on a closely watched and controlled proxy machine (or bastion) when fine-grained controls for backend services are not accessible. Access to sensitive services is restricted to requests coming from these designated proxies only. This proxy may rate limit activities, impose restrictions on harmful behaviors, and carry out more sophisticated logging.

For instance, you could have to do an urgent rollback of a problematic modification. The processes necessary to accomplish a rollback may not be included in a specified API or tool since there are an unlimited number of ways that a negative change can occur and an endless number of ways that it can be fixed. While allowing a system administrator some latitude in handling emergencies, you should also impose limitations or other risk-reduction measures. For instance:

1. Each order could need peer approval.
2. Only the proper computers may be connected to by an administrator.
3. It's possible that an administrator's machine doesn't have internet connectivity.
4. You may turn on more detailed logging.

### Tensions and Tradeoffs

Adopting a least privilege access strategy will unquestionably improve the security posture of your company. However, you must weigh the possible expense of adopting that position against the advantages described in the earlier sections. Some of such expenses are examined in this section.

- a. Complexity of Security Has Increased:** Although it is a very effective tool, a very granular security posture is both complicated and difficult to administer. To create, administer, analyse, push, and debug your security rules, you must have a complete set of tools and infrastructure. Otherwise, this intricacy could be too much to handle. These fundamental inquiries: "Does a given user have access to a given service/piece of data?" should always be your goal. and "Who has access to a given service/piece of data?"
- b. Impact on teamwork and corporate culture:** While sensitive data and services are probably best served by a rigid model of least privilege, some situations may benefit from a more permissive approach[10].

Giving software programmers extensive access to the source code, for instance, entails some risk. Engineers are allowed to learn on the job based on their own interests and may contribute features and bug fixes outside of their regular duties when they have the time and resources to do so. A less evident benefit of this openness is that it makes it more difficult for engineers to develop bad code that goes undetected.

- c. **Systems and Quality Data That Affect Security:** Every granular security decision in a zero-trust environment, which is the foundation of least privilege, relies on two factors: the policy being enforced and the context of the request. A vast amount of data, some of it possibly dynamic, informs context and may influence the choice. A training set supplied into a machine learning model, the sensitivity of the API being accessed, the properties of the client making the request, the role of the user, the groups to which the user belongs, etc. In order to guarantee that the quality of security-impacting data is as good as possible, you should assess the systems that provide this data. Incorrect security judgements will be made as a consequence of poor data quality.
- d. **User Productivity Affect:** The processes of your users must be completed as quickly as feasible. Your end users won't notice the optimal security posture, therefore adopt it. However, adding extra three-factor and multi-party permission stages might reduce user productivity, particularly if users have to wait for authorizations to be granted. By making the new processes simple to follow, you may reduce user suffering. End users also want an easy means to understand access rejections, whether it is via self-service diagnosis or quick access to a help channel.
- e. **Complexity of Developer Impact:** Developers must follow the least privilege paradigm as it is implemented across your organisation. You must make the ideas and guidelines simple to understand for developers who aren't very security-savvy, thus you must provide training materials and completely describe your APIs. Give developers quick and simple access to security experts for security evaluations and general counselling as they traverse the new specifications. In this setting, you must be very cautious when deploying third-party software since you may need to encase it in a layer that can enforce the security policy.

## CONCLUSION

The least privilege model is the most secure technique to make sure that clients can only do what they need to do while developing a complicated system. This effective design paradigm will shield your systems and data against unintentional or malicious harm caused by known or unidentified users. Google has put a lot of time and effort into putting this idea into practice. The main elements are as follows:

1. An in-depth understanding of your system's operation, allowing you to categorise individual components based on the degree of security risk they provide.
2. Based on this categorization, you should divide your system and restrict access to your data as precisely as you can. For least privilege, small, functional APIs are required.
3. An authentication process for authenticating people that try to access your system, and checking their credentials.

4. An encrypted policy-enforcing authorization system that is simple to connect to your finely partitioned systems.
5. A set of sophisticated rules for complex authorization. These may provide temporary, multiple-factor, and multi-party approval, for instance.
6. A list of necessary system operations for your system to support these important ideas. Your system must have the following, bare minimum.
7. The ability to produce signals and audit every access in order to detect risks and do forensic history analysis.
8. The resources you need to formulate, specify, test, and debug your security policy, as well as to provide end users support for it.
9. Providing a breakglass mechanism when your system behaves differently than you anticipate.

An organizational commitment to making adoption of least privilege as smooth as feasible is also necessary in order to make all these components operate in a manner that is simple for users and developers to embrace and that has little effect on their productivity. A dedicated security function that is responsible for your security posture and interacts with users and developers via security advice, policy definition, threat detection, and help on security-related problems is a part of this commitment. We firmly think that this is a big advance over current methods for enforcing security postures, despite the fact that it might be a sizable task.

## REFERENCES

- [1] S. V. Matos, M. C. Schleper, S. Gold, and J. K. Hall, "The hidden side of sustainable operations and supply chain management: unanticipated outcomes, trade-offs and tensions," *Int. J. Oper. Prod. Manag.*, 2020, doi: 10.1108/IJOPM-12-2020-833.
- [2] J. R. McColl-Kennedy, L. Cheung, and L. V. Coote, "Tensions and trade-offs in multi-actor service ecosystems," *J. Bus. Res.*, 2020, doi: 10.1016/j.jbusres.2020.06.055.
- [3] I. Oshri, S. L. Pan, and S. Newell, "Managing trade-offs and tensions between knowledge management initiatives and expertise development practices," *Manag. Learn.*, 2006, doi: 10.1177/1350507606060982.
- [4] S. S. Wong, "Distal and local group learning: Performance trade-offs and tensions," *Organ. Sci.*, 2004, doi: 10.1287/orsc.1040.0080.
- [5] D. G. Kolb, K. Dery, M. Huysman, and A. Metiu, "Connectivity in and around Organizations: Waves, tensions and trade-offs," *Organ. Stud.*, 2020, doi: 10.1177/0170840620973666.
- [6] S. A. M. Drury, D. Andre, S. Goddard, and J. Wentzel, "Assessing Deliberative Pedagogy: Using a Learning Outcomes Rubric to Assess Tradeoffs and Tensions," *J. Deliberative Democr.*, 2016, doi: 10.16997/jdd.245.
- [7] C. Boone, "Legal Empowerment of the Poor through Property Rights Reform: Tensions and Trade-offs of Land Registration and Titling in Sub-Saharan Africa," *J. Dev. Stud.*, 2019, doi: 10.1080/00220388.2018.1451633.
- [8] M. Siegner, J. Pinkse, and R. Panwar, "Managing tensions in a social enterprise: The

- complex balancing act to deliver a multi-faceted but coherent social mission,” *J. Clean. Prod.*, 2018, doi: 10.1016/j.jclepro.2017.11.076.
- [9] A. M. McGahan, M. L. A. M. Bogers, H. Chesbrough, and M. Holgersson, “Tackling Societal Challenges with Open Innovation,” *Calif. Manage. Rev.*, 2021, doi: 10.1177/0008125620973713.
- [10] D. J. Roux *et al.*, “Cultural ecosystem services as complex outcomes of people–nature interactions in protected areas,” *Ecosyst. Serv.*, 2020, doi: 10.1016/j.ecoser.2020.101111.

## DESIGN FOR UNDERSTANDABILITY

**Mr. Jayaraj Dayalan\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA,  
Email Id:dayalanj@presidencyuniversity.in

---

### ABSTRACT:

*In Site Reliability Engineering (SRE), designing for understandability entails developing procedures and systems that are simple to grasp, debug, and maintain. The significance of understandability in SRE is examined in this abstract, along with important factors to take into account. Maintaining dependable and scalable systems is crucial in the quick-paced world of technology. However, it's equally crucial to make sure that these systems are clear to the SRE teams in charge of running them. SREs can discover faults, diagnose difficulties, and make defensible judgements to restore service dependability by designing for understandability. Several elements must be taken into account in order to obtain understandability. First off, thorough documentation is essential for giving SREs the knowledge they need on the architecture, parts, and dependencies of the system. SREs are better able to comprehend system behavior and fix issues when the code, configuration files, and operational runbooks are well documented.*

**KEYWORDS:** Abstraction, Documentation, Modularity, Standardization.

---

### INTRODUCTION

Site Reliability Engineering (SRE) teams and other stakeholders may readily grasp and navigate systems, procedures, and documentation if they are designed with understandability in mind. Understanding is essential for facilitating internal cooperation and information exchange inside an organisation, as well as managing and troubleshooting complicated dispersed systems.

Here are some crucial factors to take into account while building SRE for understandability:

#### Documentation

The system design, operating methods, dependencies, and troubleshooting instructions should all be explained in clear and succinct documentation. SRE teams can immediately comprehend how various components interact and spot possible problems thanks to well-documented systems.

#### System Visibility

Utilise observability and monitoring solutions that provide thorough insights into the operation and behaviour of the system. SRE teams are better able to diagnose problems and comprehend the current condition of the system thanks to real-time dashboards, logs, and analytics [1]–[3].

**Standardization and Consistency**

Create uniform norms, naming conventions, and design patterns for all systems. Standardisation facilitates cognitive load reduction and makes it simpler for SRE teams to explore and comprehend various system components.

**Abstraction and Modularity**

Create systems with modular parts and distinct limits to contain complexity and encourage comprehension. Understanding and analysing complex systems' behaviour and interconnections is made simpler by disassembling them into manageable pieces.

**Error Handling and Alerting**

Implement useful warning systems and efficient error handling procedures. System understandability during incidents is improved through error messages and alarms that are clear and practical. This aids SRE teams in making timely diagnoses and fixing problems.

**Automation and Tooling**

Create tools and automation that streamline routine processes for SRE teams and other stakeholders and encourage self-service. Tools for automated deployment, setup, and testing lessen the mental strain brought on by manual operations and make it simpler to comprehend and operate the system.

**Onboarding and training**

Offer thorough onboarding and training programs for new SRE team members. SREs should get enough training on the system architecture, tools, and procedures to assure comprehension and promote productive cooperation.

**Collaborative Culture**

Encourage a collaborative and knowledge-sharing environment both within the SRE team and across the organisation. Regular meetings, incident reviews, and knowledge-sharing sessions promote the dissemination of information and experience, which eventually enhances system comprehension.

Organisations may improve the comprehension of their SRE practices by taking these factors into account when designing and implementing systems and procedures. This results in improved system dependability overall, quicker troubleshooting, and more effective incident management. Effective monitoring and warning systems can help make things more understandable. SREs can swiftly spot abnormalities and comprehend their influence on system performance by building monitoring systems that provide insightful and responsive notifications. SREs are assisted in identifying the main cause and swiftly implementing remedial measures by clear and informative alarm messages combined with the proper alert levels.

Additionally, adopting simplicity and eschewing needless complication are also components of designing for understandability. Complex systems and procedures might make it more difficult to grasp and lengthen the troubleshooting process. SREs may lessen cognitive burden, expedite processes, and improve their capacity to understand and address problems by putting an emphasis on simplicity. The usage of standardised, user-friendly tools and interfaces is another factor to take into account. Understanding system behavior and locating problems is made easier

by creating user-friendly command-line interfaces, dashboards, and debugging tools. By facilitating easy integration and communication with other systems, consistent and well-documented APIs also aid in making things more understandable.

As a result, sustaining dependable and scalable systems depends on planning for understandability in SRE. SREs are able to grasp system behavior, identify issues, and come to wise conclusions because of good documentation, efficient monitoring, simplicity, simple interfaces, and standard tools. By putting understandability first, SRE teams are better able to swiftly restore service dependability, reduce downtime, and improve user experience as a whole. You must be able to effectively reason about and comprehend the system, its components, and their interactions if you want to have confidence in the security posture of your system and its capacity to meet its service level objectives (SLOs). The level of a system's understandability varies greatly across distinct attributes. For instance, it could be simple to comprehend a system's behavior under heavy loads but challenging to comprehend the system's behavior when it comes into contact with maliciously designed (specially crafted) inputs.

Every step of the system lifetime is covered in this chapter's discussion of system understandability. We begin by talking about how to evaluate and comprehend your systems using invariants and mental models. We demonstrate how a layered system architecture with standardised identity, permission, and access control frameworks may assist you in designing for comprehension. We examine how programme architecture, particularly the usage of application frameworks and APIs, may dramatically influence your capacity to reason about security and dependability characteristics after delving further into the subject of security boundaries.

## **DISCUSSION**

### **Why Is Understandability Important?**

It takes work to create systems that are simple to comprehend and to keep them that way over time. As described in Chapter 4, this work often pays off in the form of sustained project velocity. An intelligible system provides the following unique advantages:

#### **Reduces the possibility of security flaws or resilience problems**

Every time you update a system or software component for instance, when you add a feature, correct a bug, or change the configuration there is an inherent risk that you might unintentionally create a new security vulnerability or jeopardise the operational resilience of the system. The engineer altering the system is more prone to make a mistake the less intelligible the system is. This engineer may not be aware of a hidden, implied, or undocumented requirement that clashes with the modification or may have the wrong understanding of how the system currently behaves.

#### **Allows for efficient incident response**

Responders must be able to evaluate damage after an event quickly and correctly, contain it, and locate and address its underlying causes. That procedure is substantially hampered by a complicated, difficult-to-understand system.

#### **Boosts the credibility of claims made regarding a system's security stance**



Typically, claims regarding a system's security are defined in terms of invariants, which are characteristics that must hold regardless of the system's conceivable behaviors. This includes the how the system responds to unforeseen interactions with its surroundings, such as when it receives incorrect or maliciously constructed inputs. In other words, a malicious input cannot cause the system to behave in a way that compromises a necessary security attribute. It may be difficult or perhaps impossible to confirm such claims with a high degree of confidence in a complex system. Testing often only exercises the system for a relatively small proportion of potential behaviors that correspond to normal or anticipated functioning, hence testing is frequently inadequate to verify that a "for all possible behaviors" quality holds. Usually, in order to establish such qualities as invariants, you must depend on abstract reasoning about the system.

### **System Invariants**

No matter how the environment of the system acts or doesn't behave, a system invariant is a property that is always true. Even if the environment of the system acts in an arbitrarily unexpected or malevolent manner, the system is entirely responsible for guaranteeing that a desired attribute is in fact an invariant. That environment consists of anything over which you have no direct control, from malevolent users who send requests to your service frontend that are specifically designed to fail to random hardware failures that cause crashes. Finding out if certain desirable qualities are in fact invariants is one of our key objectives is when analysing a system.

- a. Here are some examples of needed security and system dependability features:
- b. The persistent data store of a system is only accessible to verified and legitimate users.
- c. According to the system's auditing policy, an audit log is maintained for each operation on sensitive data in a system's persistent data store.
- d. Before being supplied to APIs that are vulnerable to injection vulnerabilities (such as SQL query APIs or APIs for creating HTML markup), any data received from beyond a system's trust boundary are adequately verified or encoded.
- e. The volume of inquiries that a system's backend receives grows in proportion to the volume of queries that the system's frontend receives.
- f. A system's frontend gracefully declines if the backend is unable to answer a query after a certain length of time, for example by providing an approximation.
- g. In order to lower the danger of cascading failure, a component will serve overload errors rather than failing when the demand is more than it can manage[4]–[6].
- h. A system is only able to send RPCs to and receive RPCs from a certain set of specified systems.

Your system has a security flaw or vulnerability if it permits actions that violate a desired security property, or if the claimed property isn't genuinely an invariant. Consider, for instance, that your system's implementation of property 1 from the list is false because an access check is either absent from a request handler or was done improperly. You now have a security hole that might let an attacker access the personal information of your users.

Consider a scenario in which your system fails to meet the fourth property: in certain cases, it creates too many backend requests for each incoming frontend request. For instance, if a backend

request fails or takes too long, the frontend can produce many retries quickly (and without a proper backoff mechanism). Your system has a possible availability weakness: if it reaches this condition, the frontend can totally overload the backend and render the service unavailable, creating a situation akin to a self-inflicted denial-of-service.

### **Analysing Invariants**

There is a trade-off when determining whether a system satisfies a certain invariant between the potential damage brought on by breaches of that invariant and the work required to satisfy the invariant and confirm that it truly holds. On one end of the scale, such work may include running a few tests and looking over some of the source code to search for errors that might cause the invariant to be violated, such as forgotten access checks. This method does not provide a very high level of confidence. Behaviour that has not been subjected to testing or a thorough code review may very well include problems. It is significant that well-known popular types of software vulnerabilities, including as buffer overflows, cross-site scripting (XSS), and SQL injection, continue to dominate "top vulnerability" lists. Lack of evidence does not prove absence.

On the other end of the spectrum, you could conduct analyses based on formal reasoning that is provably sound: you model the system and the claimed properties in a formal logic and then create a logical proof (typically with the aid of an automated proof assistant) that the property applies to the system. This strategy is challenging and labor-intensive. A demonstration of the complete correctness and security of a microkernel's implementation at the machine code level, for instance, was developed as part of one of the biggest software verification efforts to date; this project required around 20 person-years of work. Formal verification is becoming realistically useful in certain circumstances, such as the construction of complicated cryptographic libraries or microkernels, although it is often impractical for the development of large-scale software applications.

This chapter seeks to propose a useful compromise. You may get a good level of confidence in these claims and support principled (but still informal) arguments that the system has certain invariants by explicitly designing systems with the purpose of understandability. This strategy has shown to be quite useful for large-scale software development at Google and very successful in lowering the frequency of common types of vulnerabilities[7]–[9].

### **Imaginary Models (Mental Models)**

It is challenging for humans to think holistically about very complex systems. In actual practice, engineers and subject matter specialists often create mental models that describe key system behaviors while omitting unimportant aspects. You might create many mental models that complement one another for a complicated system. This allows you to abstract away the specifics of a system or subsystem's surrounding and underlying components and replace them with their corresponding mental models when considering the behaviour or invariants of that system or subsystem.

Mental models are helpful because they make understanding a complicated system simpler. Mental models are likewise constrained for the same reason. It may be difficult to anticipate a system's behaviour under unexpected situations if you rely your mental model on familiarity with how the system behaves under standard operating settings. Security and reliability engineering is

mostly concerned with analysing systems under precisely those peculiar circumstances, such as when a system is actively being attacked, when it is overloaded, or when a component fails.

Think about a system whose throughput typically rises reliably and steadily as the number of requests grows. But once the system reaches a specific load level, it could enter a state where it behaves quite differently. For instance, memory stress may cause thrashing<sup>6</sup> at the heap/garbage collector level, preventing the system from handling the increased load. A surplus of extra workload may even cause throughput to decline. If you don't expressly acknowledge that the model no longer applies while troubleshooting a system in this condition, you risk being badly misled by an excessively simplistic mental representation of the system. It is beneficial to take into account the mental models that software, security, and reliability engineers will unavoidably create for themselves while developing systems. A new component's naturally forming mental model should, in theory, be compatible with the mental models individuals have developed for related existing subsystems when it is designed to be added to a larger system.

When it is feasible, you should build systems such that even when they are put through severe or unexpected circumstances, their mental models continue to be accurate and helpful. For instance, you may set up production servers to function without on-disk virtual memory swap space in order to prevent thrashing. A production service may rapidly and predictably deliver an error if it is unable to allocate the memory it need to process a request. You may at least explicitly attribute the failure to an underlying issue in this example, memory pressure so that the mental models of the people viewing the system remain helpful, even if a flawed or misbehaving service cannot manage a memory allocation failure and crashes[10].

### **Designing Understandable Systems**

The rest of this chapter outlines some practical steps you may take to improve a system's comprehension and to preserve a system's comprehension over time. We'll start by thinking about the complexity problem.

### **Complexity vs Comprehensibility**

Complexity that isn't under control is the main threat to understandability. Because of the size of contemporary software systems especially distributed systems and the issues they attempt to tackle, some level of complexity is often inherently present and inescapable. For instance, Google has over ten thousand developers working in a source repository with over one billion lines of code. Together, those lines of code construct a sizable number of user-facing services, as well as the supporting backends and data pipelines. Even smaller businesses may utilise hundreds of thousands of lines of code, modified by tens or hundreds of engineers, to implement hundreds of features and user stories for a single product offering.

As an example of a system with substantial intrinsic feature complexity, consider Gmail. Gmail may be summed up in a single sentence as a cloud-based email service, but that basic description betrays its complexity. Gmail has the following capabilities among its many others:

- i. Multiple UIs and frontends (desktop web, mobile web, and mobile applications)
- ii. There are many APIs that let other developers create add-ons.
- iii. IMAP and POP inbound and outbound interfaces.

- iv. Integrated attachment management for cloud storage services.
- v. Rendering of attachments in a variety of formats, including spreadsheets and documents.
- vi. A web-based client with offline functionality and the underlying synchronization infrastructure.
- vii. Filtering spam.
- viii. Automatic classification of messages.
- ix. Systems for extracting structured data from calendar events, flights, etc.
- x. Spelling correction.
- xi. Smart Compose and Smart Reply.
- xii. Reminders to reply to messages.

We can't really tell Gmail's product managers that these features add too much complexity and urge them to remove it for the sake of security and reliability as a system with such features is intrinsically more complicated than a system without them. After all, the features provide value and significantly characterize Gmail as a product. But the system may still be adequately safe and dependable if we put in the effort to handle its complexity. As was already established, understandability is important when discussing certain characteristics and behaviors of systems and subsystems. Our objective must be to compartmentalize and confine this inherent complexity in a system's architecture in a manner that makes it possible for a person to reason accurately about these particular, important system traits and behaviors. In other words, we must precisely handle the complexity-related factors that obstruct comprehension.

This is obviously easier said than done. The remainder of this section looks at design patterns that may assist keep complexity under control and increase system understandability as well as frequent causes of uncontrolled complexity and its accompanying decline in understandability. The patterns we cover are mostly in line with basic software design principles aimed at controlling complexity and increasing understandability, even if security and dependability are our main concerns. A Philosophy of Software Design by John Ousterhout (Yaknyam Press, 2018), for instance, is a broad work on system and software design that you may also wish to consult.

### **Dissecting Complexity**

You need to internalize and sustain a sizable mental model if you want to comprehend every facet of a complicated system's behaviour. Simply said, humans are not very good at it. By breaking a system down into smaller parts, you may make it easier to grasp. Each component should be able to stand alone and be combined in such a manner that the qualities of the whole system may be inferred from the attributes of the individual components. With this method, you may create whole-system invariants without having to consider the system as a whole at once. In actuality, this strategy is not simple. The way the system is divided up into components, the nature of the interfaces between those components, and the relationships of trust between those components all affect your ability to define subsystem attributes and combine subsystem characteristics into system-wide properties. In System Architecture, we'll examine these connections and associated factors.

### **Centralized Management of Security and Reliability Standards**

As was said, requirements for security and dependability often apply horizontally to all parts of a system. For instance, a security requirement can specify that the system must carry out a routine activity (such as audit logging and operational metrics gathering) or verify a condition (such as authentication and authorization) before carrying out any action in response to a user request. It might be challenging to tell if the system genuinely meets the requirement if each component is responsible for individually implementing similar activities and tests. By giving a centralized component often a library or framework control over common functions, you may enhance this architecture. An RPC service framework, for instance, may guarantee that the system implements authentication, authorization, and logging in accordance with a policy that is established centrally for the whole service for each and every RPC method. These security functions aren't the responsibility of individual service methods under this approach, therefore application developers can't neglect to implement them or implement them poorly. A security reviewer may also comprehend a service's authentication and authorization restrictions without having to study every implementation of every service method. The reviewer just has to comprehend the framework and go at the settings particular to each service.

Another example would be that inbound requests should be subject to time-outs and deadlines in order to avoid cascade failures during periods of high traffic. Any logic that retries overload-related failures need to be subject to strict safety controls. You may depend on application or service code to establish sub request deadlines and properly handle errors in order to enforce these rules. An error or omission in any pertinent code in just one application might leave the system as a whole vulnerable to failure. By incorporating methods in the underlying RPC service architecture to facilitate automated deadline propagation and centralized handling of request cancellations, you may increase the stability of a system and make it easier to comprehend. Automation and tooling streamline processes and encourage self-service, lowering physical labor and mental load. Programs for training and onboarding make sure that SRE teams are familiar with the system, tools, and procedures. A collaborative culture promotes information exchange, fostering a community of knowledge and skill. Overall, improving system dependability, hastening incident response, and fostering effective internal cooperation are all benefits of designing for understandability in SRE. By putting these guidelines into practice, businesses may enable SRE teams to efficiently manage complex systems, reduce downtime, and provide users a smooth experience.

### **CONCLUSION**

For managing and troubleshooting complex systems successfully, Site Reliability Engineering (SRE) must be designed for understandability. Organisations may improve the comprehension of their systems and processes by emphasizing comprehensive documentation, system visibility, standardisation, modularity, error management, automation, training, and cooperation. SRE teams may more easily comprehend a system's design, dependencies, and operational processes through improving understandability. It lessens cognitive strain, makes effective troubleshooting possible, and encourages quicker incident response. While system visibility via monitoring and observability technologies delivers real-time insights into system behaviour, clear documentation offers a complete reference. Systems become more user-friendly and straightforward to traverse via standardisation and uniform design principles. For a better understanding of system

components and interactions, complexity is broken down via modularity and abstraction. Clear signs of problems are provided by effective error handling and alerting methods, assisting in quick diagnosis and repair.

## REFERENCES

- [1] D. Ahmed, "Anthropomorphizing artificial intelligence: towards a user-centered approach for addressing the challenges of over-automation and design understandability in smart homes," *Intell. Build. Int.*, 2021, doi: 10.1080/17508975.2020.1795612.
- [2] M. G. Al-Obeidallah, "The impact of design patterns on software maintainability and understandability: A metrics-based approach," *ICIC Express Lett. Part B Appl.*, 2021, doi: 10.24507/icicelb.12.12.1111.
- [3] V. A. Ubbes, A. M. Witter, C. M. Kraska, and E. E. Justus, "Evaluation of an eBook for Oral Health Literacy© to Promote Child Health: Readability, Suitability, Understandability, Actionability, and Gist-Based Message," *Child. Teenagers*, 2020, doi: 10.22158/ct.v3n1p54.
- [4] D. Groeneveld, B. Tekinerdogan, V. Garousi, and C. Catal, "A domain-specific language framework for farm management information systems in precision agriculture," *Precis. Agric.*, 2021, doi: 10.1007/s11119-020-09770-y.
- [5] X. Shu, A. Wu, J. Tang, B. Bach, Y. Wu, and H. Qu, "What makes a data-GIF understandable?," *IEEE Trans. Vis. Comput. Graph.*, 2021, doi: 10.1109/TVCG.2020.3030396.
- [6] D. Taibi and V. Lenarduzzi, "On the Definition of Microservice Bad Smells," *IEEE Softw.*, 2018, doi: 10.1109/MS.2018.2141031.
- [7] N. S. Mani, T. Ottosen, M. Fratta, and F. Yu, "A health literacy analysis of the consumer-oriented covid-19 information produced by ten state health departments," *J. Med. Libr. Assoc.*, 2021, doi: 10.5195/jmla.2021.1165.
- [8] A. Kurti, F. Dalipi, M. Ferati, and Z. Kastrati, "Increasing the Understandability and Explainability of Machine Learning and Artificial Intelligence Solutions: A Design Thinking Approach," in *Advances in Intelligent Systems and Computing*, 2021. doi: 10.1007/978-3-030-74009-2\_5.
- [9] C. Czepa, A. Amiri, E. Ntentos, and U. Zdun, "Modeling compliance specifications in linear temporal logic, event processing language and property specification patterns: a controlled experiment on understandability," *Softw. Syst. Model.*, 2019, doi: 10.1007/s10270-019-00721-4.
- [10] M. Jones and M. Smith, "Traditional and alternative methods of measuring the understandability of accounting narratives," *Accounting, Audit. Account. J.*, 2014, doi: 10.1108/AAAJ-04-2013-1314.

## SYSTEM ARCHITECTURE AND SOFTWARE DESIGN

Dr. Ganpathi Chandankeri\*

\*Associate Professor,  
Department Of Civil Engineering, Presidency University, Bangalore, INDIA,  
Email Id:chandankeri@presidencyuniversity.in

---

### ABSTRACT:

*Large-scale distributed systems' reliability, scalability, and maintainability are crucially dependent on the system architecture and software design used in Site Reliability Engineering (SRE). In the context of SRE, this abstract gives a general overview of the major ideas and factors to be taken into account while designing systems and software. In SRE, a system's overall structure and individual components must be designed to fulfil reliability and performance standards. To guarantee ongoing operation under a variety of circumstances, it takes into account factors including fault tolerance, load balancing, scalability, and resilience. The goal of SRE teams is to build systems that can manage rising user demand, bounce back from errors, and adjust to changing needs without sacrificing dependability. For effective development, debugging, and troubleshooting, modular, reusable, and maintainable code is the core goal of SRE software design. It entails making proper programming language, framework, and library selections and following accepted software engineering best practices. Design patterns and concepts that support flexibility, extension, and testability are prioritised by SRE teams to make maintenance and problem-solving simpler.*

**KEYWORDS:** *Configuration Management, Loose Coupling, Robust Security Parameters, System Architecture.*

---

### INTRODUCTION

In Site Reliability Engineering (SRE), system architecture and software design are essential for assuring the dependability, scalability, and maintainability of distributed systems. SRE teams concentrate on creating systems that are resilient to outages, scalable to meet rising demand, and simple to administer and maintain. SRE also emphasizes the significance of automation in software and system architecture. SRE teams may decrease human error, increase productivity, and improve system dependability by automating common processes including deployment, scaling, and monitoring. Additionally, automation speeds up problem reaction times and makes proactive maintenance and capacity planning possible.

Additionally, observability and monitoring are given priority in SRE system architecture and software design. To see system behavior, performance, and faults, SRE teams use logging, metrics, and tracing technologies. This observability enables proactive problem detection and resolving, promoting rapid recovery and reducing downtime. In the context of SRE, system architecture and software design depend heavily on collaboration and communication. To guarantee a common knowledge of needs, limitations, and tradeoffs, SRE teams collaborate closely with development, operations, and other stakeholders. Successful decision-making,

efficient problem-solving, and cross-functional knowledge exchange are all made possible through successful teamwork.

Creating dependable, scalable, and maintainable distributed systems requires careful consideration of both system architecture and software design in SRE. SRE teams may develop architectures and designs that allow effective operation, quick incident response, and continuous improvement in system reliability and performance by taking fault tolerance, scalability, automation, and cooperation into consideration. Several important factors for system architecture and software design in SRE are listed below:

### **Resilience and Redundancy**

To lessen the effects of failures, design systems with redundancy and fault-tolerant features. To guarantee high availability and resilience, this involves implementing load balancing, replication, and failover mechanisms.

### **Scalability**

Build systems that can support expansion and increasing loads. To guarantee that systems can grow effectively, this entails using horizontal scaling strategies, such as data partitioning, distributed computing, and using cloud services.

### **Loose Coupling**

In order to reduce dependencies and facilitate independent development and deployment, systems should be designed with loose coupling between components. This makes fault isolation, upgrades, and maintenance simpler [1]–[3].

### **SOA: Service-Oriented Architecture**

Adopt a service-oriented design that uses clearly defined APIs to allow for communication between various components. This encourages modularity, permits independent scalability, and makes fault management and isolation easier.

### **Observability and Monitoring**

To learn more about the behavior and performance of the system, include monitoring and observability features into the system design. This entails gathering and examining metrics, logs, and traces in order to identify bottlenecks, address problems, and enhance system performance.

### **Configuration Management**

To maintain consistency and manageability of system settings, use strong configuration management practices. To make deployment and maintenance easier, this incorporates version control, automation, and centralization of configuration settings.

### **Self-Healing and Automation**

Automate routine administrative operations and include self-healing features into the system. As a result, incident response and recovery are sped up, dependability is increased, and manual intervention is decreased.



### **Deployment and Testing**

To allow frequent and dependable software releases, use continuous integration and continuous deployment (CI/CD) practices. Automated testing, canary deployments, and progressive rollouts are all part of this strategy to lower the possibility of introducing new problems.

### **Security and Compliance**

Include security and compliance controls in the software and system architecture. To safeguard data and guarantee regulatory compliance, this requires putting into place secure coding practices, access restrictions, encryption, and auditing methods.

### **Documentation and Knowledge Sharing**

To encourage comprehension and information sharing among SRE teams and other stakeholders, document the system architecture, design choices, operating processes, and troubleshooting recommendations. SRE teams can create robust, scalable, and managed systems that can endure failures, adapt to changing needs, and allow effective operations and incident management by taking these factors into account throughout system architecture and software design.

## **DISCUSSION**

### **System Architecture**

A crucial method for controlling complexity is layering and componentizing systems. By using this strategy, you may reason about the system in segments rather than needing to comprehend every aspect of it all at once. You must carefully consider the precise division of your system into its levels and components. Too closely connected components are just as difficult to comprehend as monolithic systems. You must pay equal attention to the borders and interfaces between components as well as the components themselves in order to make a system intelligible.

A system must see inputs from (and sequences of interactions with) its external environment as unreliable, and a system cannot make assumptions about those inputs, according to experienced software engineers. The temptation is to trust callers of internal, lower-layer APIs (such as those of in-process service objects or RPCs offered by internal backend microservices) and depend on them to adhere to the API's specified use restrictions.

Let's say that a system security feature relies on an internal component functioning properly. Consider as well that its proper functioning is dependent on the API callers for the component, such as limitations on the values of method arguments or the proper order of actions. Understanding the API's implementation and every call site the API makes across the whole system as well as if each of these call sites guarantees the necessary precondition is necessary to ascertain whether the system genuinely has the desired feature. It is simpler to reason about a component independently the less presumptions it makes about its callers. A component should never assume anything about the people it calls. If a component must make assumptions about its callers, it is crucial to clearly include these assumptions via interface design or other environmental limitations, such as limiting the major types that may interact with the component.

### **Simple and Understandable Interface Specifications**

A system's comprehension is aided by structured interfaces, consistent object models, and idempotent operations. These factors, which are discussed in more detail in the next sections,

make it simpler to forecast output behaviour and interface interactions[4]–[6]. Choose interfaces with little space for interpretation. Services may offer interfaces using a wide range of models and frameworks. Among these are a few:

1. RESTful HTTP with JSON with OpenAPI
2. gRPC
3. Thrift
4. W3C Web Services (XML/WSDL/SOAP)
5. CORBA
6. DCOM

While some of these models provide more structure than others, others are quite flexible. For instance, a service that makes use of gRPC or Thrift specifies the name of each RPC method it supports in addition to the sorts of input and output the method may provide. While the application code verifies that the request body is a JSON object with the appropriate format, a free-form RESTful service may accept any HTTP request. It is simpler to develop tools for features like cross referencing and compliance checks that improve the discoverability and understandability of an API surface when using frameworks that accept user-defined types (such as gRPC, Thrift, and OpenAPI). These frameworks often enable the gradual, safer growth of an API surface. As an example, OpenAPI includes API versioning as a standard feature. There are clear instructions on how to change message definitions in protocol buffers, which are used to declare gRPC interfaces, while maintaining backward compatibility.

In contrast, unless you examine the actual code and fundamental business logic of an API based on free-form JSON strings, it might be difficult to comprehend. This unrestrained approach might result in dependability or security issues. An RPC payload, for instance, can be interpreted differently by a client and a server if they are changed separately, which might result in one of them crashing. It is challenging to assess the service's security posture in the absence of a clear API definition. For example, it would be difficult to develop an automated security audit system to correlate the rules stated in an authorization framework like Istio Authorization Policy with the actual surface area exposed by services unless you had access to the API description.

### **Pay attention to idempotent operations**

When repeated, an idempotent operation will provide the same outcome. For instance, if someone presses the floor two button in an elevator, the elevator will always proceed to the second level. Even pressing the button more than once won't affect the result. Idempotency is crucial in distributed systems because it prevents operations from arriving out of order or preventing a server's answer from ever reaching a client after an action is complete. A client may retry an operation until it is successful if an API method is idempotent. The system may need to utilise a backup strategy, such as polling the server to check if a freshly generated object already exists, if a method isn't idempotent.

Engineers' mental models are impacted by idempotence as well. Results might be unreliable or inaccurate if an API behaves differently than what is expected of it. Consider a scenario where a customer wishes to add a record to a database. Despite the request being successful, the


connection being reset prevents the delivery of the answer. The client will probably retry the request if the developers of the client code think the procedure is idempotent. However, the system will generate a duplicate record if the operation is not genuinely idempotent.

Idempotent operations often result in a simpler mental model, even when no idempotent actions may sometimes be required. Engineers (including developers and incident responders) don't have to keep track of when an operation began when it is idempotent; they may just keep attempting the operation until they are certain it works. Other operations can be made idempotent by reorganizing them, while other operations are idempotent by nature. In the previous example, the database may request that the client provide each modifying RPC with a unique identification (such as a UUID). The server may react appropriately if it detects a second mutation with the same unique identifier since it knows the action is a duplication.

### Access control

A net advantage of using frameworks to define and implement access control restrictions for incoming service requests is that the whole system will be easier to comprehend. Frameworks are a crucial component of an engineer's toolbox because they reinforce common knowledge and provide a uniform manner to articulate regulations. Frameworks are able to manage relationships that are fundamentally complicated, such as the various identities required to move data across workloads. Figure 1 illustrates the following, for instance:

1. A chain of workloads running as three identities: Ingress, Frontend, and Backend.
2. An authenticated customer making a request[7]–[9].



**Figure 1: Interactions involved in transferring data between workloads.**

The framework must be able to detect whether the workload or the customer is the request's authority for each point in the chain. In order to determine which workload identity is permitted to retrieve data on behalf of the customer, policies must be sufficiently expressive. The majority of engineers can comprehend these controls because they have a single, uniform mechanism to record this intrinsic complexity. Understanding would be difficult if each service team used its own ad hoc solution to handle the same complicated use case.

Consistency in defining and implementing declarative access control rules is required by frameworks. Engineers may create tools to assess the security exposure of services and user data inside the infrastructure thanks to this declarative and uniform nature. Creating such tools would be almost difficult if the access control logic were done in an ad hoc manner at the application code level.

### Security Limitations

A system's trusted computing base (TCB) is "the set of components (hardware, software, human,) whose correct functioning is sufficient to ensure that the security policy is enforced, or, more poignantly, whose failure could result in a breach of the security policy"<sup>10</sup> As a result,

even if a third party acts arbitrarily and perhaps maliciously, the TCB must respect the security policy. In addition to the external environment of your system (such as malevolent actors on the internet), the region outside of the TCB also contains components of your own system that are not within the TCB.

A security border is the point at which the TCB and "everything else" interact. The TCB interacts with "everything else" by exchanging information over this barrier, including other system components, the outside environment, clients of the system that communicate with it through a network, and so on. This communication may take the shape of network packets, interprocess communication channels, or higher-level protocols (like gRPC) built on top of those. Everything that crosses the security barrier must be regarded with suspicion by the TCB, including the data itself and related elements like message ordering.

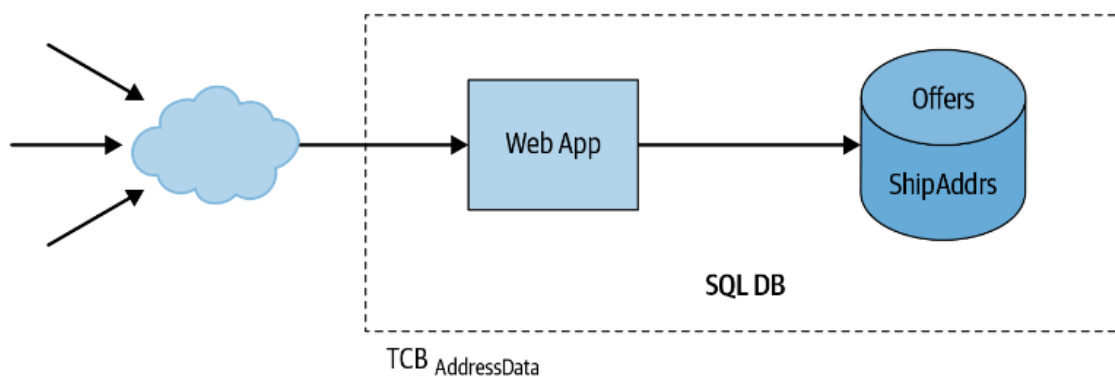
The system components that make up a TCB are determined by the security policy you have in mind. Consideration of security rules and the TCBs required to maintain them in layers might be helpful. For instance, an operating system's security model often refers to "user identity" and offers security rules that mandate isolation between operations carried out by various users. In Unix-like systems, processes owned by separate users A and B shouldn't be able to access or alter each other's memory or network traffic.<sup>11</sup> At the software level, the operating system kernel, together with all privileged processes and system daemons, make up the TCB that guarantees this property. In turn, the operating system often depends on components like virtual memory that are made available by the underlying hardware. The TCB that relates to security regulations governing user separation at the OS level includes several methods.

Since it operates as a nonprivileged OS-level role (such as the http user), the software of a network application server (for instance, the server providing a web application or API) is not covered by the TCB of this OS-level security policy. That application, however, could impose its own security policy. Consider a scenario in which a multiuser application has a security policy that restricts access to user data to explicit document-sharing controls only. If so, the application's code (or a part of it) complies with this application-level security policy inside the TCB.

You need to comprehend and rationalize the whole of the TCB pertinent to that security policy in order to guarantee that a system applies the required security policy. A flaw or error in any component of the TCB might, by definition, lead to a violation of the security policy. As a TCB expands to incorporate more code and complexity, reasoning about it gets more challenging. Because of this, it's important to maintain TCBs as manageable as possible and to get rid of any parts that aren't genuinely necessary for preserving the security policy. The inclusion of several unconnected components increases risk and reduces understandability, since a fault or failure in any of these components might lead to a security breach [10].

Let's go back to the web application that enables online widget purchases from our Chapter 4 example. Users may submit their payment card and shipping address information throughout the checkout process of the application's user interface. The system keeps some of the data and sends other portions (like credit card information) to an outside payment provider. Only the users themselves should be able to view their own private information, such as shipping addresses. For this security attribute, the trusted computing base will be indicated by `TCBAddressData`. We may

end up with an architecture similar to Figure 2 if we choose one of the several widely used application frameworks.

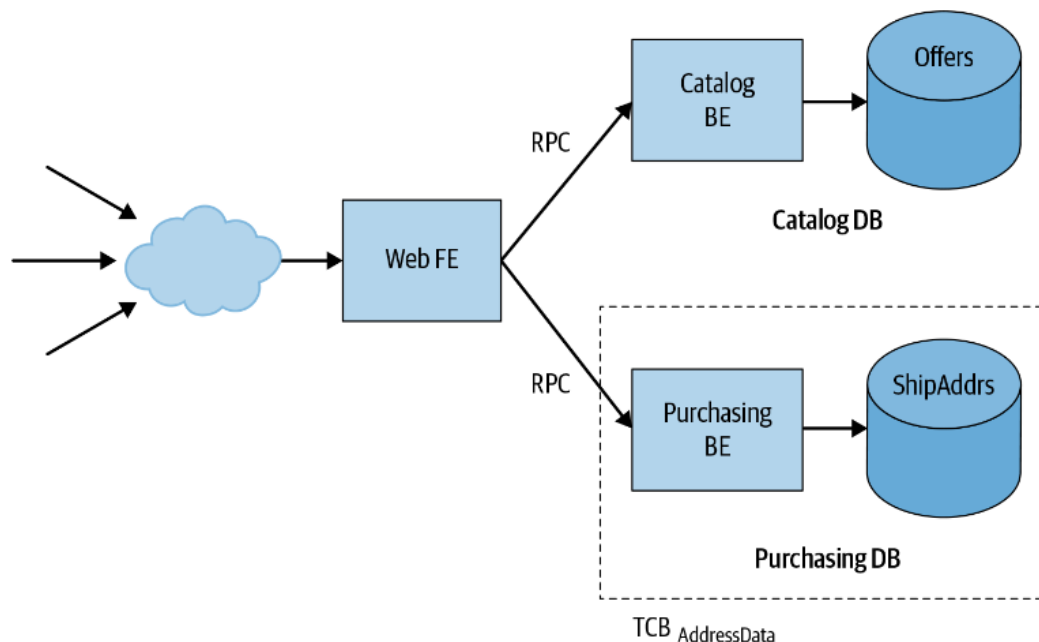


**Figure 2: Example architecture of an application that sells widgets.**

### Small TCBs and robust security parameters

By dividing the application into microservices, we may increase the security of our architecture. Each micro service in this design manages a self-contained portion of the application's operation and keeps data in a different database. These microservices use RPCs to interact with one another and disregard all incoming requests, regardless of whether they come from within microservices.

We might reorganize the application using microservices as seen in Figure 3.



**Figure 3: Example microservices architecture for widget-selling application.**

We now have a web application frontend and distinct backends for the product catalogue and purchasing-related operations instead of a single monolithic server. Each backend has a distinct

database of its own. Never directly querying a database, the web frontend instead sends RPCs to the relevant backend. To find goods in the catalogue or to get information about a specific item, for instance, the frontend queries the catalogue backend. The frontend also sends RPCs to the buying backend to conduct the checkout procedure for the shopping cart. Backend microservices and database servers may use infrastructure-level authentication protocols like ALTS to authenticate callers and restrict requests to authorized workloads, as was covered previously in this chapter.

## CONCLUSION

Systems that are easy to comprehend are deeply and intricately beneficial to both reliability and security. Although "reliability" is commonly used interchangeably with "availability," this term really refers to sustaining all of a system's crucial design guarantees, including availability, durability, and security invariants, to mention a few. Our main recommendation for creating a system that is easy to comprehend is to create it with parts that have specific, well-defined goals. Its trusted computing base may be made up of some of those components, which would concentrate accountability for resolving security risk. We also covered methods for ensuring desired characteristics inside and across those components, including security invariants, architectural resilience, and data persistence. Understanding your system's behavior when your most crucial system behaviors aren't working properly might be the difference between a quick incident and a long-lasting catastrophe. To perform their duties, SREs must be aware of the system's security invariants. In severe circumstances, they can be forced to shut down a service during a security issue, giving up availability for security.

## REFERENCES

- [1] M. Lukaszewicz *et al.*, "System architecture and software design for electric vehicles," in *Proceedings - Design Automation Conference*, 2013. doi: 10.1145/2463209.2488852.
- [2] C. Yu, Q. Li, K. Liu, Y. Chen, And H. Wei, "Industrial Design and Development Software System Architecture Based on Model-Based Systems Engineering and Cloud Computing," *Annu. Rev. Control*, 2021, doi: 10.1016/j.arcontrol.2021.04.011.
- [3] D. Spinellis and P. Avgeriou, "Evolution of the Unix System Architecture: An Exploratory Case Study," *IEEE Trans. Softw. Eng.*, 2021, doi: 10.1109/TSE.2019.2892149.
- [4] J. Tummers, A. Kassahun, and B. Tekinerdogan, "Reference architecture design for farm management information systems: a multi-case study approach," *Precis. Agric.*, 2021, doi: 10.1007/s11119-020-09728-0.
- [5] PMI, "Project Management Professional Handbook," *Pmi*, 2020.
- [6] A. Gopalakrishnan and A. C. Biswal, "Quiver-An intelligent decision support system for software architecture and design," in *Proceedings of the 2017 International Conference On Smart Technology for Smart Nation, SmartTechCon 2017*, 2018. doi: 10.1109/SmartTechCon.2017.8358574.
- [7] P. G. Teixeira *et al.*, "Constituent System Design: A Software Architecture Approach," in *Proceedings - 2020 IEEE International Conference on Software Architecture Companion, ICSA-C 2020*, 2020. doi: 10.1109/ICSA-C50368.2020.00045.

- [8] Y. Tian, D. Jing, C. Yang, Y. Chen, and H. Yang, "An ontological approach for architecture design of a smart tourism system-of-systems," *Int. J. Performability Eng.*, 2020, doi: 10.23940/ijpe.20.04.p10.587598.
- [9] European Commission, "Turkey 2020 Report," *Commun. EU Enlarg. Policy*, 2020.
- [10] S. Jones, S. Walker, M. Gatford, and T. Do, "Peeling the onion: Okapi system architecture and software design issues," *J. Doc.*, 1997, doi: 10.1108/EUM0000000007191.

---

## DESIGN FOR A CHANGING LANDSCAPE

**Mr. Narayana Gopalakrishnan\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:gopalakrishnan@presidencyuniversity.in

---

### ABSTRACT:

*It is essential for long-term success to build systems that can adapt in the quickly changing technology environment of today. This abstract examines the idea of "Design for a Changing Landscape" and its relevance to software engineering, product development, and corporate strategy, among other fields. Anticipating and accepting changes in technology, market trends, customer wants, and regulatory requirements are key components of designing for a changing environment. It necessitates adopting a proactive mentality that works to include resilience, scalability, and flexibility into the very foundation of systems and processes.*

**KEYWORDS:** *Containerization, Deployment, Flexibility, Integration, Scalability.*

---

### INTRODUCTION

To guarantee that systems can adapt to changing technologies, business demands, and user needs, Site Reliability Engineering (SRE) must take into account the changing environment. Systems must be designed to be flexible, robust, and simple to adapt to in order to keep up with the quickly evolving technical and commercial world. Here are some important factors to keep in mind while building systems for a changing environment:

**Decoupled and modular architecture:** Create systems with modular parts that can be changed or upgraded separately. When changes are performed, cascading failures may be avoided thanks to loose coupling between the components.

**APIs and microservices:** Adopt an API-defined microservices architecture. As a result, various services may develop separately and be quickly merged or changed as required, enabling flexibility and agility[1]–[3].

**Continuous Integration and Deployment (CI/CD):** To allow quick and frequent software releases, use CI/CD practices. By making it easier to deliver new features, bug patches, and security upgrades, systems are better able to keep up with evolving needs.

**Automation:** Accept automation for mundane processes like infrastructure provisioning, scalability, and monitoring. Automation lessens the need for physical labor and makes it possible to respond to changing needs more quickly.

**Elasticity and Scalability:** System designs should grow horizontally to accommodate rising loads. To ensure systems can manage variations, use cloud services and auto-scaling features to automatically modify resources depending on demand.



**Observability and Monitoring:**To learn more about the behavior and performance of your system, use reliable monitoring and observability practices. Early problem detection is made possible by proactive monitoring, which also enables timely optimizations.

**Security and Compliance:**Integrate compliance and security into the system architecture. Keep abreast of industry best practices, routinely evaluate risks, and put security policies in place to safeguard systems and data in a dynamic threat environment.

**Iteration and Experimentation:**Encourage an environment of experimentation and ongoing development. To make sure systems can adapt to changing user demands and preferences, encourage teams to test new concepts, get feedback, and iterate on designs.

**Documentation and Knowledge Sharing:**Decisions about the system's architecture and operating processes should be documented. To establish a common understanding and enable future revisions, promote knowledge exchange across teams and maintain current documentation.

**Collaboration and Cross-functional Teams:**Encourage cooperation across various groups, such as the development, operations, and business stakeholder teams. Cross-functional teams may cooperate to comprehend shifting requirements and create solutions that address changing demands.

SRE teams may build flexible and resilient systems that can survive in a changing environment by taking these factors into account. Flexibility, automation, continuous improvement, and collaboration are crucial foundations for navigating the possibilities and difficulties brought on by a corporate environment and technology that are always evolving. Designing for a changing terrain should take into account important variables like:

#### **Iteration and agility**

Adopt agile processes and iterative strategies to react swiftly to new possibilities and difficulties. Create procedures and systems that are simple to modify and enhance in response to user input and shifting needs.

#### **Interoperability and Modularity**

Create systems with modular parts that may be updated or replaced separately. In order to facilitate smooth integration with external systems or changing industry requirements, promote interoperability across various components.

#### **Scalability and Performance**

Create systems that can scale effectively to meet rising needs and workloads. To dynamically modify resources depending on demand variations, take into account employing cloud-based services and elastic infrastructures.

#### **User-Centric Design**

To inform design choices, collect user input and insights consistently. Prioritize user pleasure and experience, and make sure that systems can change to accommodate changing user demands and preferences.

#### **Future-Proofing**

Be prepared for future changes in business and technology breakthroughs. Make systems extensible, compatible, and adaptable to make it simpler to integrate new technologies or business models.

### **Data-Driven Decision Making**

Utilise data insights and analytics to guide design decisions. Analyse data continuously to recognise patterns, identify possible threats or opportunities, and direct iterative changes.

### **Continuous Learning and Skill Development**

Develop a culture of skill development and ongoing learning inside the organisation. Encourage teams to keep current on market trends, test out new technology, and learn what they need to in order to adjust to a changing environment. Designing for a changing environment include business models, strategies, and organizational structures in addition to technology. In order to overcome uncertainty and take advantage of new possibilities, a comprehensive strategy that integrates technology, people, and processes is needed.

## **DISCUSSION**

The adage "Change is the only constant"<sup>1</sup> is undoubtedly true for software; as the number (and variety) of devices we use grows yearly, so do the vulnerabilities in libraries and applications. Any software or hardware might be vulnerable to remote exploit, data leaking, botnet takeover, or other in-the-news events. Simultaneously, consumers' and regulators' demands for security and privacy are increasing, necessitating tougher controls like enterprise-specific access limits and authentication systems. You must be able to update your infrastructure often and fast while still providing a highly dependable service if you want to be able to adapt to this ever-changing world of vulnerabilities, expectations, and dangers. This is no small task. Choosing when and how soon to implement a change is often the key to achieving this equilibrium.

### **Types of Security Changes**

You might implement a variety of different improvements to raise your security posture or the resilience of your security infrastructure, such as:

- a. Changes in response to security incidents
- b. Changes in response to newly discovered vulnerabilities
- c. Changes to products or features
- d. Internally driven adjustments to strengthen your security posture
- e. Changes brought on by outside forces, including new legal requirements

Certain security-related modifications need more thought. In order to get enough input from early adopters and adequately test your first instrumentation, you must first roll out a feature on an optional basis before making it essential. Make sure the new solution satisfies your security needs before changing a dependence, such as one that depends on a vendor or third-party code.

### **Architecture Decisions to Make Changes Easier**

How can your infrastructure and procedures be designed to adapt to the unavoidable changes you'll experience? Here, we go through various tactics that help you adapt your system and

implement changes with the least amount of resistance while simultaneously fostering a culture of security and dependability.

### **Keep Dependencies Up to Date and Rebuild Frequently**

Your system is less prone to new vulnerabilities if your code always refers to the most recent versions of its dependencies. For open source projects that change often, such as OpenSSL or the Linux kernel, updating references to dependencies is crucial. When a new release includes a crucial security patch, it is made clear in many big open source projects' defined security vulnerability response and remediation strategies, and the repair is back ported to supported versions. Instead of having to merge with a backlog of changes or apply numerous patches, if your dependencies are current, you probably can apply a key fix straight[4]–[6].

Until you rebuild, new releases and associated security updates won't reach your environment. In order to be prepared to roll out a new version when necessary and to ensure that an emergency rollout can include the most recent changes you should often rebuild and reload your environment.

### **Release Frequently Using Automated Testing**

To support emergency adjustments, basic SRE principles include cutting and rolling out releases often. You may make sure that each release has fewer changes and is thus less likely to need rollback by dividing a single major release into several smaller ones. It's simpler to comprehend what changed and identify possible problems when each release involves fewer code changes. You may have greater faith in the result when you need to put out a security update.

Automate testing and validation of frequent releases to fully benefit from them. As a result, effective releases may be automatically pushed while defective releases are kept out of production. Automated testing also increases your confidence when it comes time to provide updates that counter serious flaws. Similar to this, you may decrease the surface area you need to patch, set up frequent release procedures, and make it easier to comprehend system vulnerabilities by leveraging containers and microservices.

### **Use Containers**

The binaries and libraries required by your application are separated from the host OS by containers. The host OS may be smaller since each programme comes packed with all of its own dependencies and libraries. Applications become more portable as a consequence, and you may secure them separately. For instance, without altering your application container, you may repair a kernel vulnerability in the host operating system. In order for containers to remain unchanged after deployment, you must rebuild and reinstall the whole image rather than SSHing into a computer. Containers are often repaired and redeployed because of their limited lifespan.

You update the pictures in your container registry rather than active containers. The patch rollout process will now be the same as your (very frequent) code rollout process, replete with monitoring, canarying, and testing. This implies that you can deploy a completely patched container image as a single unit. You may thus patch more often. Containers may also be used to find and fix recently disclosed vulnerabilities. Containers enable content addressability since they are immutable. In other words, you are aware of what is really operating in your environment, such as the deployed pictures. Instead of physically scanning your production

clusters, you may utilise your registry to find the vulnerable versions and apply updates if you already released a completely patched image that just so happens to be vulnerable to a new vulnerability.

You should keep an eye on how old the production-level containers are running and reload often enough to make sure that they aren't. This will lessen the need for this type of ad hoc patching. Similarly, you should mandate that only freshly produced containers may be deployed in production in order to prevent redeploying older, unpatched images.

### **Use Microservices**

An optimal system design enables for simple scaling, insight into system performance, and management of any possible bottlenecks between services in your infrastructure. Workloads may be divided into smaller, easier-to-manage units using a microservices architecture to make maintenance and discovery easier. As a consequence, each microservice can independently scale, load balance, and do rollouts, giving you greater freedom to make infrastructure changes. You may deploy a number of defences individually and in succession to provide defence in depth since each service handles requests differently (see Defence in Depth). On addition, restricted or zero trust networking is naturally facilitated by microservices, which means that your system does not automatically trust a service simply because it is present on the same network. Microservices employ a more diverse sense of trust within the perimeter: internal traffic may have varied degrees of trust, as opposed to a perimeter-based security model with trusted external vs untrusted internal traffic. The direction of current developments is towards a more divided network. The network may be further divided into services when reliance on a single network perimeter, such a firewall, is reduced. In its most extreme form, a network may segregate services at the microservice level without any inherent trust between them. The convergence of security technologies as a result of the usage of microservices allows certain procedures, tools, and dependencies to be shared throughout several teams. It could make sense to combine your efforts to solve similar security needs as your design expands, for instance by adopting a common infrastructure for monitoring and alerting or cryptography libraries. By doing this, you may divide important security services into discrete microservices that are updated and controlled by a small group of accountable individuals. In order to preserve the appropriate security features while keeping the services as basic as possible, restraint is necessary to achieve the security benefits of microservices architectures.

Teams may handle security vulnerabilities in a standardised manner early in the development and deployment lifecycle, when it's less expensive to make changes. This is made possible by using a microservices architecture and development methodology. Developers may provide safe results while devoting less effort to security as a consequence [7]–[9]. By employing a microservices architecture, constructing many levels of security controls, and using a service mesh to manage cross-service interactions in the modern cloud environment, you may obtain advantages like to those mentioned above. As an instance, you might divide request processing from the setup for handling requests. This kind of intentional division is referred to in the business as "separation of the data plane" (the requests) and "control plane" (the configuration). In this paradigm, the data plane handles load balancing, security, and observability as well as the actual data processing inside the system. A manageable and scalable control surface is provided by the control plane by way of the policy and configuration it applies to the data plane services.

### **Various Changes: Varying Speeds and Timelines**

Not every change takes place in the same timeframe or at the same rate. What determines how rapidly you may wish to change depends on a number of factors:

#### **Severity**

Every day, new vulnerabilities are found, but not all of them are serious, being actively exploited, or relevant to your specific architecture. You probably want to deploy a fix as soon as possible when you do achieve that trifecta. The likelihood of systems breaking increases with accelerated timescales. While speed is sometimes required, it is normally preferable for a change to develop gradually so that you can guarantee adequate product security and dependability. (Ideally, you should be able to deploy a critical security patch on your own. This will allow you to complete the fix swiftly without needlessly speeding other in-flight rollouts.)

#### **Teams and dependent systems**

Some system modifications can be reliant on other teams that have to roll out new procedures or activate a specific feature first. Your modification could also be reliant on a third party, such as if you need a patch from a vendor or if clients must be fixed prior to servers.

#### **Sensitivity**

When you can deploy your modification to production may depend on how sensitive it is. An improvement to an organization's overall security posture via a non-essential modification may not need the same level of urgency as a critical patch. You may introduce that unnecessary change more gradually, perhaps team by team. Making the change may not be worth the risk depending on other variables; for instance, you might not want to implement a non-urgent change during crucial production periods like the Christmas shopping season, when modifications are often strictly restricted.

#### **Deadline**

Certain updates have a set time limit. For instance, a regulation modification can have a deadline for compliance, or you might need to install a patch before a news embargo (see the sidebar below) revealing a vulnerability becomes public. A modification that needs a rapid configuration update and deployment in one organisation may take months in another. There is no hard-and-fast rule for assessing the pace of a specific change. While one team may be able to implement a change in a particular timeframe, it could take your organisation some time to completely absorb the change [10].

We examine three alternative time horizons for change in the sections below, and we use examples from Google to illustrate each.

- i. A quick adjustment in response to a fresh security flaw
- ii. A medium-term shift where acceptance of new products could occur gradually
- iii. A long-term regulatory shift that required Google to develop new processes in order to implement

### **Short-Term Change: Zero-Day Vulnerability**

Newly identified vulnerabilities often call for immediate response. A zero-day vulnerability is one that hasn't been publicly reported or identified by the infrastructure provider being attacked, but is at least partially understood by attackers. Usually, a patch is either not yet accessible or hasn't been extensively used. Regular code reviews, internal code scanning (see Sanitize Your Code), fuzzing (see Fuzz Testing), external scans like penetration testing and infrastructure scans, and bug bounty programs are just a few methods to learn about new vulnerabilities that might impact your environment.

We'll concentrate on issues where Google was made aware of the vulnerability on day zero in the context of short-term adjustments. Although Google often participates in embargoed vulnerability responses, such as when creating patches, a quick fix for a zero-day vulnerability is typical practise for the majority of businesses in the sector. Determine the impact and severity of a new vulnerability when it is found. A major vulnerability may be one that permits remote code execution, for instance. But it could be difficult to gauge the effect on your company: Which systems use this specific binary? Is the impacted version active in real-world settings? To ascertain if the vulnerability is being actively exploited, you'll also want to set up regular monitoring and alerts when it is practical.

To take action, you must get a patch, which is a fresh copy of the problematic library or package with the repair applied. Check to see whether the patch genuinely fixes the vulnerability first. Using a functional exploit to do this may be beneficial. Be cautious that your system may still be susceptible even if you are unable to execute the exploit to trigger the vulnerability (remember that a lack of evidence does not prove a presence). For instance, the patch you installed could only cover one potential attack among a group of vulnerabilities.

Once your fix has been validated, release it ideally in a test environment. The same testing, canarying, and other methods should be used to roll out a fix gradually, on the order of hours or days, even on an expedited timeframe. As the patch could have an unanticipated impact on your applications, a slow rollout enables you to identify possible problems early. For instance, an application that uses an API that you were not aware of might affect performance or result in other issues.

Sometimes the vulnerability cannot be immediately fixed. In this situation, blocking or otherwise restricting access to the susceptible components is the recommended line of action for risk mitigation. This mitigation may be permanent if you are unable to install the patch to your system, for example due to performance constraints, or temporary until a patch is available. You may not even need to take any more action if your environment is already secured by adequate mitigations.

### **Medium-Term Change: Improvement to Security Posture**

Changes are often made by security teams to strengthen an environment's overall security posture and lower risk. Rarely do these proactive adjustments need to be implemented unexpectedly since they are driven by needs and deadlines from both internal and external sources. You need to identify which teams and systems are impacted when making changes to your security posture and choose the best location to start. Designing Your Change's SRE principles should be followed while you create an action plan for a phased deployment. There should be success requirements for each step that must be satisfied in order to go to the next.

Security modifications to systems or teams may not always be reported as a proportion of a deployment. Instead, you may roll out in phases depending on who will be impacted and what modifications need to be made. Roll out your change incrementally, taking into account the individuals who will be impacted. A group may be a system, a development team, or a collection of end users. For instance, you may start by implementing a modification to the device settings for users that travel a lot, like your sales staff. By doing this, you can rapidly test the most typical instances and get feedback from actual users. When it comes to rollout populations, there are two contrasting philosophies:

- i. To get the greatest traction and demonstrate your value, start with the simplest use case.
- ii. Start with the most challenging use case since it contains the majority of defects and edge situations.

Start with the simplest use case while you're still trying to get the organisation on board. Finding implementation problems and pain points early on is more helpful if you have leadership backing and commitment up front. You should think about which method will result in the largest risk reduction, both in the short and long term, in addition to organizational considerations. A successful proof of concept is always helpful in determining the best course of action. The change-making team should experience it firsthand, "eating their own dog food," in order to fully grasp the user experience.

Additionally, you may be able to implement the modification itself gradually. You may be able to impose gradually stricter criteria, for instance, or make the change initially opt-in rather than required. Before converting to an enforcement mode, you should, wherever feasible, try rolling out a change as a dry run in an alerting or auditing mode so that users can see how they will be impacted. Using this, you may identify persons or systems that you may have mistakenly included in your scope as well as those for whom attaining compliance will be especially challenging.


### **Long-Term Change: External Demand**

In other circumstances, you either have or need a lot longer time to implement a change—for instance, a change that is internally driven and requires major architectural or system modifications, or a more extensive regulatory change that affects the whole sector. The implementation of these improvements might take many years and may be influenced or constrained by deadlines or requirements from other sources.

When undertaking a significant, multiyear project, it's important to define your objective precisely and track your progress towards it. In order to preserve consistency and make sure you take into account the relevant design considerations (see *Designing Your Change*), documentation is very important. Today's change-makers can quit the firm and need to transfer their efforts to someone else. For continued leadership support, it's critical to keep documents current with the most recent strategy and progress.

Create the necessary dashboarding and instruments to track continuing development. A configuration check or test should ideally be able to measure a change automatically, cutting off the requirement for a person to be involved. For systems impacted by the change, you should aim for compliance check coverage in the same way that you strive for large test coverage for code in your infrastructure. This instrumentation should be self-serve, enabling teams to execute both the

modification and the instrumentation, in order to grow this coverage efficiently. Transparently tracking these outcomes makes corporate reporting and communications easier while also encouraging users. You should also utilise this one source of truth for executive communications rather than duplicating labor.



**Figure 1: Percentage of HTTPS browsing time on Chrome by platform.**

It is challenging to implement any significant, protracted change inside an organisation while keeping ongoing support from the leadership. The people implementing these changes must continue to be motivated if they are to maintain momentum over time. Teams may complete the marathon by setting concrete objectives, monitoring their progress, and providing compelling instances of their influence. The long tail of implementation will always exist, so choose a plan that fits your circumstances the best. Achieving 80% or 90% adoption may have a demonstrable influence on decreasing security risk and can thus be regarded a success if a change is not necessary (by legislation, or for other reasons).

## CONCLUSION

It is very vital to distinguish between various security modifications so that impacted teams are aware of what is required of them and how much help you can provide. Take a deep breath and make a strategy the next time you're asked to make a security modification to your infrastructure. Find volunteers who are willing to test the change or begin small. Make the update self-service and employ a feedback loop to learn what isn't working for the users. Don't worry if your plans change, but also try not to be shocked. A layered approach results in few and well-managed exterior surface areas, while design choices like frequent rollouts, containerization, and microservices make both proactive upgrades and emergency mitigation simpler. A healthy system is maintained through thoughtful design and continuing documentation, both of which are



done with an eye towards change. These actions also reduce burden for your team and, as you'll see in the next chapter, increase resilience.

## REFERENCES

- [1] H. Van Dyck and E. Matthysen, "Habitat fragmentation and insect flight: A changing 'design' in a changing landscape?," *Trends in Ecology and Evolution*. 1999. doi: 10.1016/S0169-5347(99)01610-9.
- [2] E. B.-N. Sanders and P. Jan Stappers, "Co-creation and the new landscapes of design," in *Design: Critical and Primary Sources*, 2017. doi: 10.5040/9781474282932.0011.
- [3] E. Baumfeld Andre, R. Reynolds, P. Caubel, L. Azoulay, and N. A. Dreyer, "Trial designs using real-world data: The changing landscape of the regulatory approval process," *Pharmacoepidemiology and Drug Safety*. 2020. doi: 10.1002/pds.4932.
- [4] J. Frith, "Forum design and the changing landscape of crowd-sourced help information," *Commun. Des. Q.*, 2017, doi: 10.1145/3068698.3068700.
- [5] S. Mazumdar and S. Mazumdar, "Planning, design, and religion: America's changing urban landscape," *J. Archit. Plann. Res.*, 2013.
- [6] R. Cooper, "Design Research: Past, Present and Future," *Des. J.*, 2017, doi: 10.1080/14606925.2017.1257259.
- [7] A. Harrison and L. Hutton, *Design for the Changing Educational Landscape*. 2013. doi: 10.4324/9780203762653.
- [8] G. Muratovski, "Paradigm Shift: Report on the New Role of Design in Business and Society," *She Ji*, 2015, doi: 10.1016/j.sheji.2015.11.002.
- [9] A. Harrison and L. Hutton, *Design for the changing educational landscape: Space, place and the future of learning*. 2013. doi: 10.4324/9780203762653.
- [10] N. Mayer-Hamblett *et al.*, "Building global development strategies for cf therapeutics during a transitional cftr modulator era," *Journal of Cystic Fibrosis*. 2020. doi: 10.1016/j.jcf.2020.05.011.

## DESIGN FOR RESILIENCE

Dr. Jagdish Godihal\*

\*Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA,  
Email Id:drjagdishgodihal@presidencyuniversity.in

---

### ABSTRACT:

*The capacity to defend against assaults and to endure uncommon events that stress your system and influence its dependability is a key component of good system design. You should consider how to maintain the system fully or partly operational while dealing with several simultaneous events early in the design process. This chapter begins with a tale from antiquity, in which layered defence may have prevented the fall of an empire. After that, we examine contemporary defense-in-depth techniques using Google App Engine as an example. The options discussed in this chapter range in terms of implementation costs and organizational size suitability. We advise concentrating on controlled deterioration, setting up blast radius restrictions, and segmenting systems into distinct failure zones if your organisation is smaller. We advise employing continuous validation to check and improve the robustness of your system as your organisation expands. The term "resilience" refers to a system's capacity to withstand a significant breakdown or interruption as part of its design. Systems that are resilient may automatically recover from failures that affect just a portion of the system or even the whole system and resume regular operations after the issues have been fixed. In a resilient system, services should ideally continue to operate notwithstanding an event, maybe in a degraded condition. Every layer of a system's architecture should have resilience to protect it from unforeseen failures and attack situations.*

**KEYWORDS:** *Chaos Engineering, Load Balancing, Redundancy, Replication.*

---

### INTRODUCTION

Site Reliability Engineering (SRE) places a high priority on designing resilient systems that can endure disturbances, breakdowns, and unanticipated occurrences. Systems that are resilient can adjust, carry on, and lessen how events affect user experience. Here are some essential factors for creating resilient systems:

**Replication and Redundancy:** Use replication and redundancy techniques to make sure that crucial parts or services have backups. Multiple instances of components may need to be deployed in various availability zones or data centers to achieve this [1]–[3].

**Fault Isolation:** Create systems with distinct boundaries and definite component-to-component interactions. This aids in containing failures and stops them from propagating across the system.

**Elasticity and load balancing:** To uniformly distribute traffic and manage rising demands, use load balancing methods and scalable infrastructure. The system's elasticity enables it to

autonomously scale resources up or down in response to demand, assuring peak performance and reducing interruptions.

**Disaster Recovery and Failover:**To automatically transition to backup systems or services in the case of a breakdown, implement failover mechanisms. Create and test disaster recovery strategies as well in order to recover from significant catastrophes or disasters.

**Alerting and Monitoring:**To find problems and probable breakdowns, use effective monitoring and warning systems. In order to spot abnormalities and take proactive action, keep an eye on system metrics, logs, and health checks.

**Automated Incident Response:**Automated incident response techniques should be used to swiftly identify and address issues. This may include incident response runbooks, automated alerting, and automated recovery processes.

**Graceful Degradation and Circuit Breakers:**When systems are stressed or fail, they should gracefully lose performance. Use circuit breakers to separate defective components and stop failures from cascading.

**Testing and Chaos Engineering:**To proactively find system flaws and failure sites, do chaos engineering experiments and resilience testing. This increases the system's resilience while revealing possible weaknesses.

**Documentation and Knowledge Sharing:**Keep documents current and engage in knowledge exchange. This makes sure that during events, the SRE team and other stakeholders will have quick access to information on the system design, failure scenarios, and recovery techniques.

**Learning and Continuous Improvement:**Encourage an environment where learning from mistakes is valued. Update procedures and documentation, conduct post-incident evaluations to identify areas for improvement, and put preventative measures into place to boost system resilience.

SRE teams may create resilient architectures that can sustain failures, recover fast, and guarantee a high degree of availability and reliability by include these factors in the design and implementation of systems. Designing for resilience is a continuous process that calls for constant observation, evaluation, and development in order to adjust to shifting circumstances and developing threats.

## DISCUSSION

### Resilience Design Principles

The design ideas outlined previously in Part II serve as the foundation for a system's resilience characteristics. You need to have a solid grasp of the system's structure and architecture in order to assess its resilience. To improve the stability and resilience of your system, you must closely align with the other design principles discussed in this book, including least privilege, understandability, adaptability, and recovery.

1. A resilient system is defined by the following strategies, each of which this chapter discusses in detail:
2. Design the system such that each layer is robust on its own. With each layer, this strategy adds depth to the defence.

3. Consider the importance of each feature and its cost to determine which functions should be prioritized and attempted to be maintained regardless of the system's load, and which features should be throttled or deactivated if issues develop or resources are limited. Then, you may decide how to best use the system's limited resources and how to increase its capacity for providing.
4. To encourage the independence of the separated functional sections, divide the system into distinct compartments with distinct boundaries. Additionally, it is simpler to develop supplementary defence behaviors in this manner.
5. To protect against localized failures, use compartment redundancy. Have some compartments provide various levels of security and dependability in the event of worldwide failures.
6. By securely automating as many of your resilience measures as you can, you can shorten system response time. Investigate fresh failure modes that might profit from brand-new automation or enhanced automation already in place.
7. Maintain the system's efficacy by confirming its resilience capabilities, including both its automatic reaction and any other resilience-related characteristics.

### **Defence in Depth (Division)**

Defence in depth creates many levels of defence perimeters to safeguard systems. Attackers have less access to the systems as a consequence, making it more difficult to execute effective vulnerabilities.

### **The Trojan Horse**

The Trojan horse myth, as recorded by Virgil in the Aeneid, serves as a warning on the perils of a weak defence. The Greek army builds a giant wooden horse and gives it to the Trojans as a gift after ten futile years of besieging Troy. When the horse is brought close to Trojan's defenses, attackers concealed inside of it suddenly emerge, take advantage of the city's defenses from the inside, and then let the whole Greek army in through the city's gates, destroying Troy[4]–[6]. If the city had a comprehensive defence strategy, consider how this narrative might have ended. First, it's possible that Troy's defenses would have examined the Trojan horse more carefully and identified the trick. If the attackers had been able to enter the city gates, they may have encountered an additional line of defence, such as the horse being confined to a safe courtyard with no access to the rest of the city.

What can we learn about scaled security, or perhaps security itself, from a 3,000-year-old tale? To begin with, you must comprehend the assault itself in order to comprehend the tactics you need to protect and confine a system. If we think of the city of Troy as a system, we may go through the processes used by the invaders (the phases of the assault) to find areas where defence in depth might be able to help. We can categorise the Trojan assault into four phases on a high level:

- a. Assess the target and especially seek for defenses and vulnerabilities. Threat modelling. The city gates were locked from the outside, but could they be opened from the inside by the attackers?

- b. Establish the circumstances for the assault during deployment. Troy ultimately brought a thing that the enemy had built and delivered within its city gates.
- c. Execution-Execute the real assault, building on the steps taken before. Soldiers exited the Trojan horse and unlocked the city's gates so the Greek army could enter.
- d. Compromise-After the attack is carried out successfully, damage is done, and the mitigation process starts.

Before the compromise, the Trojans had several chances to halt the invasion, but they chose not to take advantage of them, and they paid a steep price. The same is true for your system's defence in depth, which may lower the cost you could incur if it is ever hacked.

### **Modelling threats and identifying vulnerabilities**

Both attackers and defenders are able to identify a target's vulnerabilities. Attackers conduct reconnaissance on their targets to identify vulnerabilities, after which they plan assaults. Defenders should take all reasonable measures to restrict the data that is made available to attackers during reconnaissance. Defenders must identify this reconnaissance and utilise it as a signal since they can't entirely stop it. Because of enquiries from outsiders regarding how the gates were guarded, the Trojan Horse's defenders may have been on high alert. They would have been more cautious when they discovered a big wooden horse at the city entrance in light of such suspicious activities.

It is danger intelligence collecting to take notice of these strangers' enquiries. You may opt to outsource some of these tasks for your own systems, and there are several methods to go about it. You might, for instance, carry out the following:

- a. Keep an eye out for port and application scans on your machine.
- b. Keep track of DNS registrations for URLs that resemble yours, since an attacker could use them for spear phishing assaults.
- c. Info about danger intelligence to buy[7]–[9].

Create a threat intelligence team to research and observe inactively the actions of known and potential threats to your infrastructure. While we don't advise small businesses to spend money in this strategy, it could pay off as your business expands. You can do a more thorough evaluation than the attacker can since you are a defender and are familiar with the workings of your system. This is crucial: you can more effectively guard against system flaws if you are aware of them. And the more you comprehend the techniques that attackers are now using or are able to exploit, the greater this influence becomes. Developing blind spots to attack vectors you deem implausible or unimportant should be avoided, it is advised.

### **Deployment of the attack**

If you are aware that an attacker is using your system for reconnaissance, you must make every attempt to find them and halt the assault. Imagine if the Trojans had decided that since the wooden horse was made by someone they did not trust, they would not allow it to pass through the city gates. Instead, they might have examined the Trojan Horse carefully before letting it inside, or they could have simply burned it on fire. In the current day, methods like network traffic inspection, virus identification, programme execution control, protected sandboxes<sup>1</sup>, and

adequate privilege granting for signalling unusual usage may all be used to identify possible threats.

### Execution of the attack

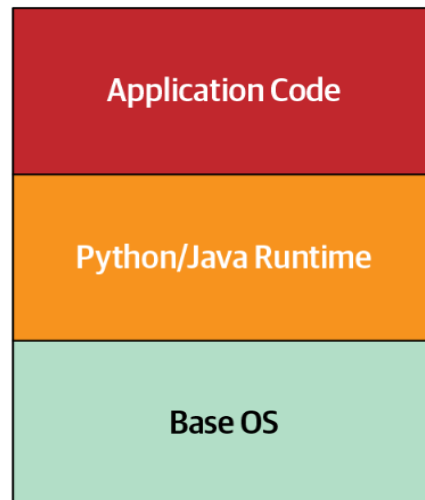
Limit the explosion radius of prospective strikes if you can't stop every deployment by your enemies. The assailants would have had a far tougher difficulty sneaking out of their hiding place unobserved if the defenders had boxed in the Trojan Horse, restricting their exposure. Sandboxing is the term used in cyberwarfare to describe this approach, which is further explained in Runtime layering.

### Compromise

The Trojans realized their city had been breached when they awoke to discover their adversaries hovering over their beds. This realization happened a long time after the actual compromise. After EternalBlue and WannaCry contaminated their infrastructure in 2018, many unlucky institutions encountered a similar predicament. How you react moving forward will decide how long your infrastructure is vulnerable.

### Analysis of Google App Engine

Let's look at defence in detail with a more recent scenario in mind: Google App Engine. Users may host application code on Google App Engine and grow as traffic rises without having to deal with networks, computers, and operating systems. An early architectural design for App Engine is shown in Figure 1. Developers are responsible for protecting the application code, while Google is in charge of protecting the Python/Java runtime and the basic OS[10].



**Figure 1: A simplified view of Google App Engine architecture.**

Special process isolation concerns were needed for Google App Engine's first implementation. We determined that executing each user's code in a separate virtual machine was too inefficient for the degree of expected adoption at the time when Google employed the typical POSIX user isolation technique (via unique user processes). We had to find out how to execute untrusted, third-party code on Google's infrastructure just like any other task.

## Risky APIs

Initially, App Engine's threat modelling identified the following concerning areas: There were issues with network access. Until that time, it was thought that any application using the Google production network was a dependable, authorized infrastructure component. We required a plan to keep internal APIs and network exposure separate from App Engine since we were adding arbitrary, untrusted third-party code into this setting. Additionally, we had to keep in mind that App Engine itself was reliant on access to those same APIs since it was operating on the same infrastructure.

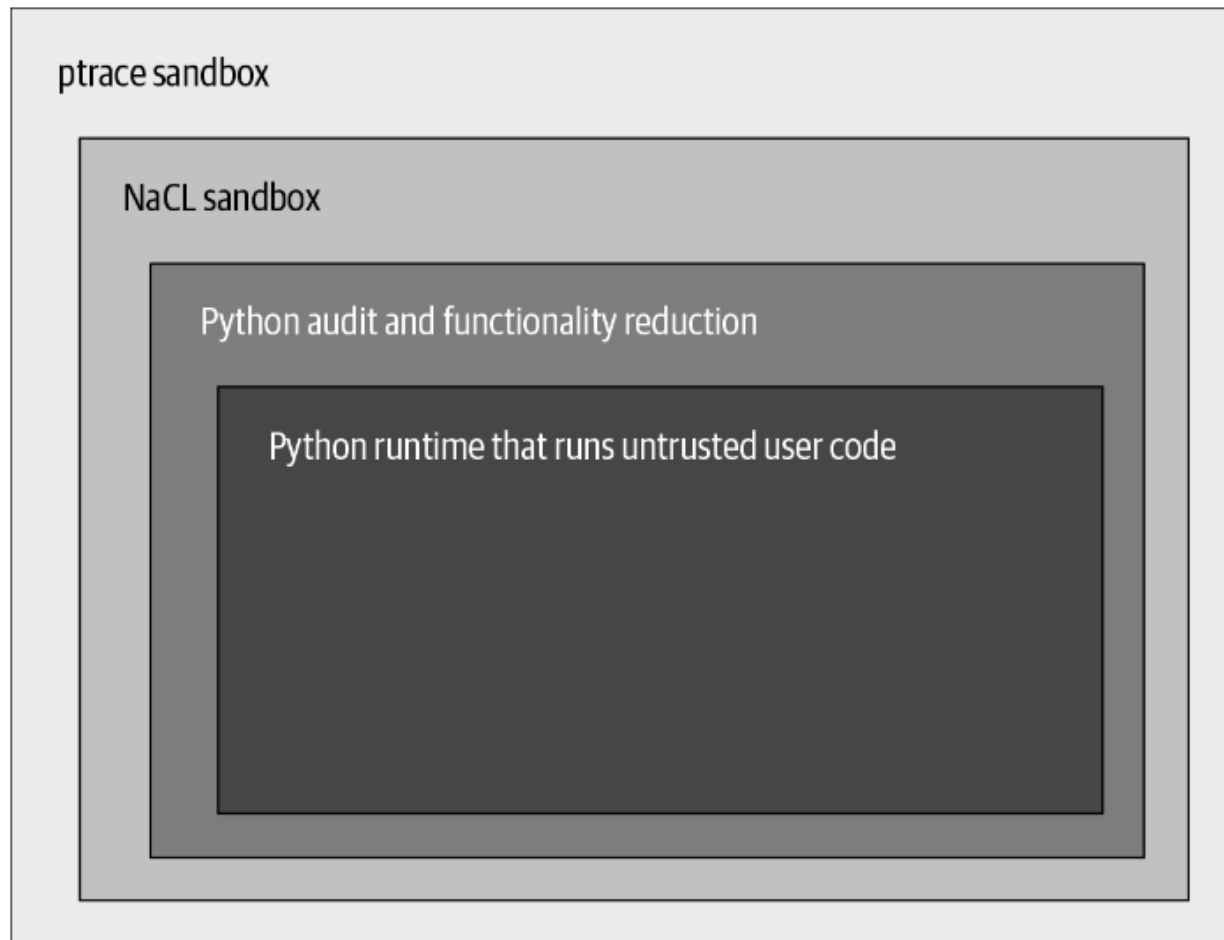
The local filesystem has to be accessible to the computers executing user programs. The execution environment was protected by the fact that this access was restricted to the directories owned by the specific user, lowering the possibility of user-provided apps interfering with those of other users on the same computer. Because of the Linux kernel, App Engine had a wide attack surface, which we intended to reduce. For instance, we intended to limit the number of classes of local privilege escalation that may occur.

We originally considered restricting user access to each API in order to solve these issues. Built-in APIs for networking and file system interfaces were deleted by our team at runtime. Instead of directly altering the runtime environment, we replaced the built-in APIs with "safe" ones that made calls to other cloud infrastructure. We prohibited user-supplied compiled bytecode or shared libraries to stop users from restoring the purposefully removed features to the interpreters. Users had to rely on our offered libraries and techniques in addition to a number of possible open source runtime-only implementations.

## Runtime layers

Additionally, we thoroughly examined runtime base data object implementations for any aspects that can lead to memory corruption problems. Several upstream bugs were fixed as a result of this audit in each of the runtime environments we released. Given that we weren't expected to identify and foresee every exploitable flaw in the selected runtimes, we predicted that at least some of these defensive measures would be unsuccessful. We particularly chose to modify the Python runtime so that it may be compiled down to Native Client (NaCL) bit code. Our thorough code auditing and hardening missed numerous kinds of memory corruption and control-flow subversion attacks, but NaCL was able to stop them.

In order to filter and warn on unusual system calls and arguments, we implemented a second layer of ptrace sandboxing to NaCl since we weren't quite certain that it would totally contain any unsafe code breakouts and flaws. Any breaches of these standards resulted in the runtime being stopped immediately, notifications being sent out with a high priority, and records of pertinent activities. The team discovered a few instances of aberrant behavior caused by exploitable circumstances in one of the runtimes during the course of the next five years. We were able to defeat the attackers in each instance (whom we later identified as security researchers) thanks to our sandbox layer, and our several levels of sandboxing kept their actions inside the intended boundaries. Functionally, the sandboxing levels shown in Figure 2 were present in App Engine's Python implementation.



**Figure 2: Sandboxing layers of Python implementation in App Engine.**

The levels of App Engine are complimentary, with each layer foreseeing the shortcomings or potential failures of the one before it. Signals of a compromise become stronger as defence activations go through the levels, enabling us to concentrate our efforts on likely attackers. Even though we used many layers of protection for Google App Engine, we still benefitted from outside assistance in protecting the environment. Our team not only found abnormal behavior, but also multiple instances of exploitable vectors were found by other researchers. We are appreciative to the researchers who identified and reported the gaps.

## CONCLUSION

Site Reliability Engineering (SRE) teams must prioritize ease of modification when making architectural choices. Businesses may reduce the work and risk involved in changing their infrastructure and applications by designing systems with flexibility, modularity, and adaptability in mind. In a dynamic technological environment where requirements, technologies, and user demands are continually changing, the capacity to make changes quickly is essential. Organisations can swiftly adapt to market needs, embrace new technologies, and constantly enhance their systems by choosing design that prioritizes ease of modification. Changes may be



made quickly, which improves the system's resilience and agility while enabling SRE teams to provide users with higher dependability and performance.

## REFERENCES

- [1] A. Kusiak, "Open manufacturing: a design-for-resilience approach," *Int. J. Prod. Res.*, 2020, doi: 10.1080/00207543.2020.1770894.
- [2] M. Drgoicea *et al.*, "Service design for resilience: A multi-contextual modeling perspective," *IEEE Access*. 2020. doi: 10.1109/ACCESS.2020.3029320.
- [3] A. Alblas and J. Jayaram, "Design resilience in the fuzzy front end (FFE) context: An empirical examination," *Int. J. Prod. Res.*, 2015, doi: 10.1080/00207543.2014.899718.
- [4] D. Liu, R. Deters, and W. J. Zhang, "Architectural design for resilience," *Enterp. Inf. Syst.*, 2010, doi: 10.1080/17517570903067751.
- [5] R. Rajesh, "Network design for resilience in supply chains using novel crazy elitist TLBO," *Neural Comput. Appl.*, 2020, doi: 10.1007/s00521-019-04260-3.
- [6] M. F. Braga, E. R. Filho, R. M. L. O. Mendonca, L. G. R. de Oliveira, and H. G. G. Pereira, "Design for resilience: Mapping the needs of Brazilian communities to tackle COVID-19 challenges," *Strateg. Des. Res. J.*, 2020, doi: 10.4013/sdrj.2020.133.07.
- [7] J. S. Baek, A. Meroni, and E. Manzini, "A socio-technical approach to design for community resilience: A framework for analysis and design goal forming," *Des. Stud.*, 2015, doi: 10.1016/j.destud.2015.06.004.
- [8] M. J. Sáenz, E. Revilla, and B. Acero, "Aligning supply chain design for boosting resilience," *Bus. Horiz.*, 2018, doi: 10.1016/j.bushor.2018.01.009.
- [9] S. Bagchi *et al.*, "Vision paper: Grand challenges in resilience: Autonomous system resilience through design and runtime measures," *IEEE Open J. Comput. Soc.*, 2020, doi: 10.1109/OJCS.2020.3006807.
- [10] M. Turnquist and E. Vugrin, "Design for resilience in infrastructure distribution networks," *Environmentalist*, 2013, doi: 10.1007/s10669-012-9428-z.

## CONTROLLING DEGRADATION AND BLAST RADIUS

**Dr. Mohammad Shahid Gulgundi\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:mohammadshahid@presidencyuniversity.in

---

### ABSTRACT:

*A crucial component of Site Reliability Engineering (SRE), which seeks to lessen the effects of failures or interruptions in distributed systems, is controlling degradation and blast radius. While explosion radius describes the degree to which a failure or event might spread and harm other system components or services, degradation describes the progressive loss in system performance. This paper examines methods and recommended procedures for limiting blast radius and degradation in SRE. It emphasizes how crucial proactive monitoring, fault isolation, and efficient incident response are in reducing the effects of failures and preventing failures from cascading across the system. To reduce the explosion radius and swiftly resume normal operation, the abstract emphasizes the necessity for robust system design, fault tolerance measures, and automated recovery methods.*

**KEYWORDS:** *Blast Radius, Degradation, Fault Isolation, Shedding, Throttling.*

---

### INTRODUCTION

For large distributed systems, controlling degradation and blast radius is a critical component of site reliability engineering (SRE), which aims to reduce the effect of failures and avoid widespread disruptions. While blast radius refers to the size of the effect a breakdown or event has, degradation relates to the decline in system performance. To maintain system reliability and availability, SRE teams work to develop and put into practice solutions that efficiently control deterioration and reduce the blast radius[1]–[3].

Implementing strategies to limit and isolate failures as well as proactively monitoring and reacting to performance concerns are necessary for controlling degradation and blast radius. By doing this, businesses may lessen the effects of failures, stop cascade failures, and improve customer experience. This allows SRE teams to detect and address faults immediately, minimizing downtime and guaranteeing the system's general stability. The paper also covers the need of using load balancing, traffic shaping, and canary deployments to limit deterioration by progressively introducing changes and controlling traffic to minimise interruptions. It also emphasizes how crucial it is to prioritize system stability and manage deterioration by establishing explicit service level goals (SLOs), using service level indicators (SLIs), and using error budgets.

To discover and mitigate vulnerabilities, better system design, and increase overall system dependability, the abstract emphasizes the necessity for a collaborative culture, constant learning, and post-incident evaluations. Organisations may assure the availability, performance, and resilience of their systems, providing a better user experience and reducing the impact of failures,

by putting policies to limit degradation and blast radius into place. In this regard, this study will examine different methods and recommended procedures for limiting blast radius and deterioration in distributed systems. It will go through tactics including monitoring, incident response, failover solutions, load balancing, chaos engineering, and fault isolation. Organisations may maintain system stability, reduce customer impact, and handle issues when they arise by putting these approaches into practice. Overall, limiting blast radius and degradation is crucial for preserving the dependability and availability of complex systems. Organisations can prevent failures, lessen their effects, and guarantee optimum system performance by putting strong plans and practices in place.

## DISCUSSION

We consider the possibility of system failure while planning for defence in depth. Failures may occur for a variety of causes, such as physical harm, hardware or network issues, software bugs or misconfigurations, or security breaches. Every system that relies on a component might be negatively impacted when it malfunctions. The total pool of comparable resources shrinks as well; for instance, system-wide storage capacity is decreased by disc failures, bandwidth and latency are increased by network failures, and compute capacity is decreased by software failures. Problems might accumulate; for instance, a storage deficit could result in software problems.

System overload might result from resource limitations like these, a rapid increase in incoming requests like those brought on by the Slashdot effect, configuration errors, or denial-of-service attacks, or both. A system's reaction always starts to deteriorate as its load exceeds its capacity, which might result in a totally unusable system. You can't predict where a system could malfunction unless you've prepared for it beforehand, but it will almost certainly happen where the system is least safe and most vulnerable.

When severe conditions develop, you must decide which system characteristics to deactivate or modify in order to limit deterioration, while still doing all in your power to preserve the system's security. For situations like this, it helps to plan for various possible responses so that instead of a chaotic collapse, the system may employ controlled breakpoints. Your system may react by gently deteriorating rather than causing cascade failures and coping with the chaos that results. Here are some strategies for doing that:

- a. By turning off seldom used features, the least important functions, or expensive service capabilities, you may free up resources and lower the rate of unsuccessful operations. The resources may then be used to maintain crucial features and functions. For instance, the majority of TLS-accepting computers offer both the RSA and Elliptic Curve (ECC) cryptosystems. One of the two will cost less while still providing you with equivalent security, depending on how your system is implemented. ECC requires less resources to operate on private keys in software. When systems are resource-constrained, disabling RSA functionality will allow for more connections at the cheaper cost of ECC.
- b. Ensure that system reaction actions happen instantly and without human intervention. With servers that are directly under your control, where you have complete control over operational parameters of any size or granularity, this is made simplest. User clients are more difficult to manage since they have lengthy rollout cycles due to the possibility that client

devices would delay or fail to receive updates. The variety of client platforms also raises the possibility of response measures being rolled back because of unforeseen incompatibilities.

- c. Recognise the systems that are essential to achieving the goals of your business, as well as their relative significance and interdependencies. Depending on how important they are, you may need to keep these systems' most basic functions. For instance, Google's Gmail offers a "simple HTML mode" that turns off sophisticated UI style and search auto completion while still allowing users to access their inboxes. Even this mode may be deprioritized if regional network problems that restrict bandwidth permitted network security monitoring to continue protecting user data.

If the system's ability to handle load or failure is significantly improved by these changes, they serve as a crucial complement to all other resilience methods and offer incident responders more time to act. Making the important and tough decisions in advance is preferable than doing so during an emergency. It is simpler to prioritize deterioration across more systems or product categories if individual systems have established a defined degradation strategy[4]–[6].

### Differentiate Failure Costs

Every unsuccessful action has a cost; for instance, a failure data upload from a mobile device to an application backend uses network bandwidth and CPU resources to build up an RPC and transfer some data. You may be able to lessen or eliminate certain failure-related waste if you can rework your processes to fail quickly or affordably.

To evaluate the cost of failures:

- a. **Determine the overall expenses of each procedure:** For instance, you may gather CPU, memory, or bandwidth impact measurements while evaluating the load of a certain API. If you're short on time, concentrate first on the most important operations—either in terms of criticality or frequency.
- b. **Find out when these expenses are incurred throughout the process:** To get introspection data, you might examine the source code or utilise developer tools (for instance, web browsers provide monitoring of request phases). Even better, you could instrument the code by include failure scenarios at various phases.

With the knowledge you get about operation costs and failure spots, you may search for adjustments that might delay more expensive operations until the system advances farther down the path to success.

### Computing Resources

Any other operations cannot use the computational resources that a failed operation uses from the start of the operation until failure. If clients actively retry after failing, this impact compounds and might possibly result in a cascading system failure. By checking for fault conditions earlier in the execution flows for instance, by verifying the legitimacy of data access requests before the system allots memory or starts data reads/writes you may release computational resources more rapidly. You may prevent dedicating RAM to TCP connection requests coming from forged IP addresses by using SYN cookies. The costliest procedures may be protected against automated misuse with the aid of CAPTCHA.

More generally, you may have a server move into a lame-duck mode<sup>4</sup> where it continues to serve but notifies its callers to throttle down or cease delivering requests if it can detect that its health is deteriorating (for instance, via the signals of a monitoring system). This method improves the signals that the ecosystem as a whole can adjust to while minimizing resources devoted to serving faults. Additionally, due to outside reasons, numerous instances of a server may be unutilized. For instance, a security breach can lead the services they manage to be "drained" or isolated. If you keep an eye out for certain circumstances, the server resources could be momentarily made available to other services. However, you should be sure to safeguard any data that can be useful for a forensic investigation before you reallocate resources.

### **User Experience**

In degraded circumstances, the system's interactions with the user should behave in an appropriate manner. An ideal solution alerts users of potential service issues while still allowing them to engage with working components. To maintain the functioning state, systems may test various connection, authentication, and authorization protocols or endpoints. Users should be made fully aware of any data staleness issues or security threats resulting from malfunctions. Explicitly disabling features that are no longer safe to use is advised.

For instance, including an offline mode in an online collaboration tool may maintain key capabilities even when online storage, the ability to display updates from others, or chat capability are temporarily lost. Users may periodically modify the encryption key that is used to secure messages in a chat application with end-to-end encryption. Because this modification has no impact on the legitimacy of prior messages, such a programme would maintain access to all of them.

On the other hand, a scenario where the whole GUI becomes unusable because one of its RPCs to a backend has run out would be an illustration of bad design. Imagine a mobile app that connects to its backends immediately after launch to show just the most recent content. Users would not even be able to access data that had previously been cached if the backends were inaccessible merely because the device's owner had purposefully stopped connection. To develop a user experience (UX) solution that offers usability and productivity in a reduced mode, a user experience (UX) research and design effort may be necessary.

### **Speed of Mitigation**

The cost of a system failure is influenced by how quickly it recovers after a failure. This reaction time takes into account the interval between a mitigating adjustment made by a person or an automated system and the update and recovery of the component's final impacted instance. Keep important sites of failure away from client apps and other harder-to-control components. Recalling the preceding illustration of the mobile application that launches a freshness update, that design decision makes access to the backends an essential need. The delayed and unpredictable Deploy response mechanisms.

A system should, in theory, actively react to worsening circumstances by taking safe, pre-programmed actions that maximise the efficacy of the reaction while minimizing risks to security and dependability. Humans are often slower to react, may not have the requisite network or security access to finish a required activity, and aren't as adept at solving for numerous variables

as automated techniques. Humans should still be involved, however, to serve as checks and balances and to make choices in unexpected or complicated situations.

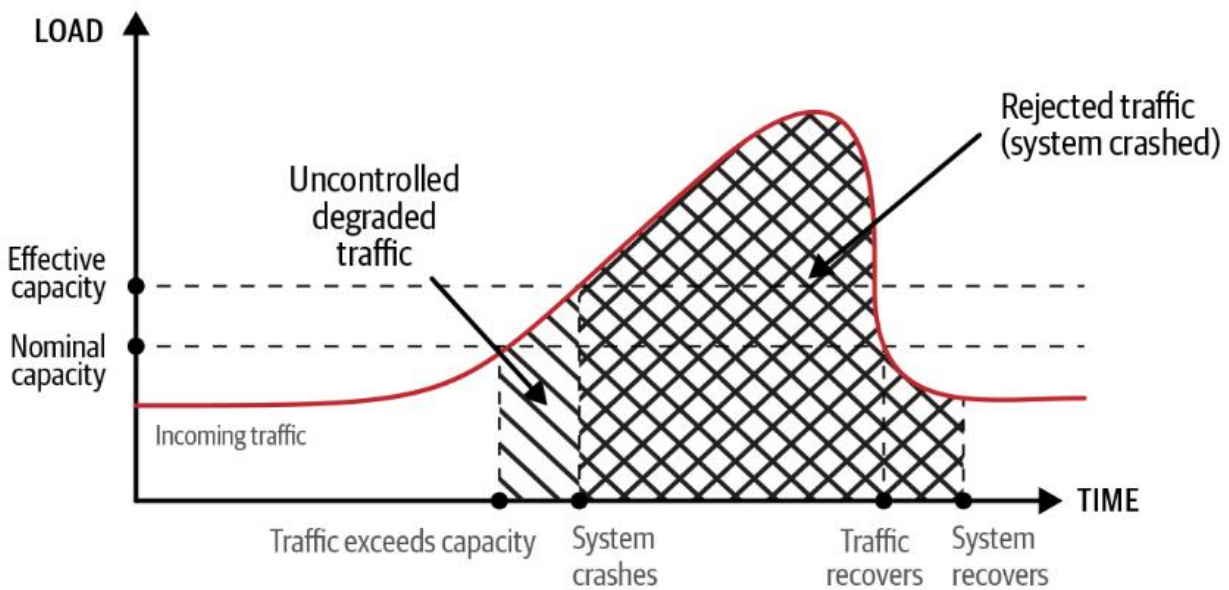
Let's talk about handling heavy loads, whether they result from DoS assaults, benign traffic surges, or a lack of serving capacity. It's possible that people won't react quickly enough, and that traffic may overload servers to the point where failures cascade and a worldwide service meltdown occurs. Permanently overprovisioning servers to serve as a safety net is a money-wasting practice that doesn't provide a secure response. Instead, servers should modify how they react to demand depending on the circumstances at hand. Here, you may use two distinct automation techniques:

- a. Instead of fulfilling requests, load shedding is accomplished by returning errors[7]–[9].
- b. By postponing replies until nearer the request deadline, clients are throttled.


A traffic surge that surpasses the capacity is seen in Figure 1. The consequences of load shedding and throttling to handle the load surge are shown in Figure 2. Keep this in mind:

- a. The area beneath the curve reflects total requests, while the curve itself represents requests per second.
- b. Whitespace denotes successfully processed traffic.
- c. Degraded traffic (some requests failed) is represented by the region that has been backlashed.
- d. The regions with crosshatching show denied traffic (all requests were unsuccessful).
- e. Prioritized traffic is seen in the forward-slashed region (important requests were granted).

Figure 1 illustrates how the system can really collapse, having a higher effect on time and volume (number of requests missed) than anticipated due to the length of the outage. The backward-slashed region in Figure 1 illustrates the uncontrolled nature of degraded traffic just before the system crashed. The system with load shedding rejects far less traffic than in Figure 1 (the crosshatched region), with the remaining traffic either being handled without issue (the whitespace area) or being rejected if it is of lesser priority (the forward-slashed area).pace of programme updates in this case exacerbates the original issues.



**Figure 1: Complete outage and a possible cascading failure from a load spike**



**Figure 2: Using load shedding and throttling to manage a load spike.**

### Load shedding

Stabilizing components at maximum load is the main resilience goal of load shedding, which might be particularly useful for maintaining security-critical operations. You want a component to deliver errors for all excessive requests rather than crashing when the load on it begins to surpass its capacity. Not only the capacity for the extra requests, but the whole component's capacity becomes unavailable in the event of a crash. The load simply moves to another location when this capacity is exhausted, perhaps leading to a cascade failure.

Even before a server's demand hits capacity, load shedding enables you to free up server resources and make them available for more useful work. The server requires an understanding of request priority and cost in order to decide which requests to drop. You may provide a policy that decides how many requests of each category to reject based on the priority, cost, and server load at the time. Prioritize requests depending on their dependencies or the request's business criticality (security-critical functions should be given high priority). Request costs may be measured or empirically estimated. Regardless, these statistics ought to be equivalent to those used to gauge server utilization, such as CPU and (potentially) memory consumption. Of course, it should be cost-effective to calculate request expenses [10].

### **Throttling**

By delaying the current operation to delay future operations, throttling indirectly alters the client's behavior. The server may wait after receiving a request before processing it, or it may wait after processing a request before giving the response to the client. If clients make requests sequentially, this method lowers the pace at which the server receives them, allowing you to repurpose the resources saved during wait periods. You might create rules to apply throttling to certain misbehaving customers, or more broadly to all clients, similarly to load shedding. The decision of which requests to throttle is influenced by request priority and cost.

### **How to Control the Blast Radius**

By restricting the scope of each component of your system, you may add another layer to your defense-in-depth approach. Think of network segmentation as an example. A single network including all of an organization's resources (computers, printers, storage, databases, etc.) was typical in the past. Any user or service on the network could access these resources, and each one had access restrictions. Segmenting your network and granting access to each section to certain types of users and services is a popular practice nowadays for enhancing security. Virtual LANs (VLANs), an easy-to-configure, industry-standard method, may be used to do this. You may manage the flow of traffic into each segment and the segments that are permitted to speak with one another. Additionally, you may restrict access to each segment's "need to know" information.

Network segmentation is a nice illustration of the compartmentalization concept in general, which we covered in Chapter 6. In order to compartmentalize, discrete, distinct operating units (compartments) are purposefully created, with access to and from each one being restricted. The majority of your system components, including servers, programs, storage, and so on, should be divided into separate areas. When you set up a single network, an attacker who steals a user's login information may be able to access any device on the network. But when you compartmentalize, a security breach or traffic jam in one compartment does not endanger the rest of them.

In the same way that compartments on a ship provide resistance against the entire ship sinking, controlling the blast radius entails compartmentalizing the effect of an event. Create compartmental barriers that impede both attackers and unintentional failures while designing for resilience. You can more effectively automate and adapt your replies thanks to these obstacles. As mentioned in Failure Domains and Redundancies, you may also utilize these boundaries to build failure domains that provide component redundancy and failure isolation. Additionally, compartments help with quarantine operations by minimizing the requirement for responders to actively balance defence and evidence preservation. While other compartments are retrieved,



certain compartments may be segregated and frozen for study. Additionally, during incident response, compartments provide a natural boundary for replacement and repair; one compartment may be destroyed in order to preserve the system as a whole.

You need a mechanism to set limits and a way to ensure that those borders are safe if you want to regulate the explosion radius of an invasion. Think of a production-running work as having one compartment. This task must allow some access in order for the compartment to be functional, but it cannot allow unlimited access since the compartment has to be protected. Your capacity to distinguish between endpoints in production and validate their identities is a prerequisite for limiting who may access the task.

Authenticated remote procedure calls, which identify both participants inside a single connection, may be used for this. These RPCs employ mutually authenticated connections, which may verify the identities of both parties connecting to the service, to secure the parties' identities from spoofing and to hide their contents from the network. You may include extra information that endpoints publish along with their identification to help them make better judgements about other compartments. For instance, you may include geographical data in the certificate so that you can deny requests from beyond your area.

Once the mechanisms for constructing compartments are in place, you are faced with a challenging trade-off: you must limit your actions with just enough separation to produce compartments of a reasonable size while avoiding over-separation. Consider each RPC method as a distinct compartment as an example of a balanced compartmentalization strategy. As a result, compartments are aligned along logical application boundaries, and the ratio of compartments to system features is linear.

More careful thought should be given to compartment separation, which regulates the permissible parameter values for RPC procedures. Although greater security measures would be put in place as a result, the amount of potential violations per RPC method is inversely correlated with the number of RPC clients. All of the system's features would become more complicated as a result, necessitating coordination between client code and server policy updates. Compartments that enclose a complete server, independent of its RPC services or their protocols, are more simpler to administer but provide far less benefit in comparison. The incident management and operations teams should be consulted while weighing the pros and disadvantages of this tradeoff so that you can confirm the usefulness of your compartment type selections.

Additionally valuable are imperfect compartments that don't completely address all edge scenarios. For instance, the search for edge situations could lead an attacker to make a mistake that lets you know they're there. Your incident response team has more time to respond the longer it takes for such an attacker to exit a compartment.

To confine an intrusion or a bad actor, incident management teams must prepare and practise their compartment closing techniques. It's a big move to turn off a portion of your production environment. The use of well-designed compartments allows incident management teams to take steps that are appropriate for the occurrences without necessarily having to shut down the whole system. When you implement compartmentalization, you must decide whether to operate distinct service instances that serve different customers or subsets of customers or have all customers share a single instance of a particular service.

For instance, there is a risk associated with operating two virtual machines (VMs) on the same hardware that are each controlled by distinct, mutually suspicious organisations. This risk may include vulnerability to zero-day vulnerabilities in the virtualization layer or subtle cross-VM information breaches. Customers may decide to compartmentalize their deployments depending on physical hardware in order to reduce these risks. Many cloud service providers provide deployment on customer-specific dedicated hardware to make this strategy easier.<sup>9</sup> In this instance, a price premium reflects the cost of lower resource utilization.

As long as the system has safeguards in place to keep the compartments separate, compartment separation increases system resilience. Monitoring such systems and making sure they stay in place is a challenging endeavor. It is important to verify that activities that are forbidden across separation borders really fail in order to avoid regressions (see Continuous Validation). Conveniently, your validation techniques may cover both banned and anticipated activities since operational redundancy depends on compartmentalization (discussed in Failure Domains and Redundancies).

Google separates work by job, place, and time. The potential reach of any one assault is significantly diminished when an attacker attempts to penetrate a compartmentalized system. The incident management teams have the ability to deactivate just a portion of the system in order to eliminate the impact of the hack while keeping the other portions operating.

## CONCLUSION

To sum up, managing degradation and blast radius is a crucial component of site reliability engineering (SRE), which aims to reduce the effects of failures and guarantee the stability of systems as a whole. Organisations may maintain service availability and minimise possible interruptions by putting procedures in place to manage the extent and effects of failures.

The following are some salient lessons on limiting deterioration and blast radius:

- a. Lowering the blast radius
- b. Fault isolation implementation
- c. Prioritizing and Load Shedding
- d. Rate Limiting and Throttling
- e. Progressive rollouts and Canary Deployments
- f. Engineering and testing amidst chaos
- g. Observation and Alerting
- h. Response to Incidents and Post-Incident Evaluation

Organisations may efficiently manage deterioration and blast radius by putting certain practices in place, which reduces the impact of failures and interruptions. System stability and dependability are increased as a result, which ultimately enhances user satisfaction.

## REFERENCES

- [1] D. Gagliardi and A. Dziembowski, "50 and 30 modifications controlling RNA degradation: From safeguards to executioners," *Philos. Trans. R. Soc. B Biol. Sci.*, 2018,

- doi: 10.1098/rstb.2018.0160.
- [2] C. Shasteen and Y. Bin Choy, "Controlling degradation rate of poly(lactic acid) for its biomedical applications," *Biomedical Engineering Letters*. 2011. doi: 10.1007/s13534-011-0025-8.
- [3] J. Yeom, Y. Shao, and E. A. Groisman, "Small proteins regulate Salmonella survival inside macrophages by controlling degradation of a magnesium transporter," *Proc. Natl. Acad. Sci. U. S. A.*, 2020, doi: 10.1073/PNAS.2006116117.
- [4] Y. A. Lee *et al.*, "Autophagy is a gatekeeper of hepatic differentiation and carcinogenesis by controlling the degradation of Yap," *Nat. Commun.*, 2018, doi: 10.1038/s41467-018-07338-z.
- [5] G. K. Sims and A. M. Cupples, "Factors controlling degradation of pesticides in soil," *Pestic. Sci.*, 1999, doi: 10.1002/(sici)1096-9063(199905)55:5<598::aid-ps962>3.3.co;2-e.
- [6] P. Nonsuwan, A. Matsugami, F. Hayashi, S. H. Hyon, and K. Matsumura, "Controlling the degradation of an oxidized dextran-based hydrogel independent of the mechanical properties," *Carbohydr. Polym.*, 2019, doi: 10.1016/j.carbpol.2018.09.081.
- [7] C. A. Woolnough, L. H. Yee, T. S. Charlton, and L. J. R. Foster, "A Tuneable Switch for Controlling Environmental Degradation of Bioplastics: Addition of Isothiazolinone to Polyhydroxyalkanoates," *PLoS One*, 2013, doi: 10.1371/journal.pone.0075817.
- [8] International Atomic Energy Agency., "Controlling of degradation effects in radiation processing of polymers (IAEA-TECDOC-1617)," *Ind. Appl. Chem. Sect. IAEA*, 2009.
- [9] M. M. Ghobashy *et al.*, "Controlling radiation degradation of a CMC solution to optimize the swelling of acrylic acid hydrogel as water and fertilizer carriers," *Polym. Adv. Technol.*, 2021, doi: 10.1002/pat.5105.
- [10] H. Shahgholi and A. G. Ahangar, "Factors controlling degradation of pesticides in the soil environment : A Review," *Agric. Sci. Dev.*, 2014.

## FAILURE DOMAINS AND REDUNDANCIES

**Mr. Ahamed Sharif\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA,  
Email Id:ahamedsharif@presidencyuniversity.in

---

### ABSTRACT:

*Redundancies and failure domains are essential ideas in Site Reliability Engineering (SRE), which aims to improve system resilience and lessen the effects of failures. Redundancies entail building backups or duplicate components to lessen the effects of failures, while a failure domain is a portion of a system or infrastructure where failures may happen separately. The importance of failure domains and redundancies in SRE is examined in this abstract, which also highlights their function in preserving system availability, dependability, and fault tolerance. The implementation of failure domains and redundancies is also covered, along with numerous tactics and best practices, such as geographic redundancy, data replication, load balancing, and failover methods. Organisations may improve their capacity to tolerate disruptions, recover rapidly, and provide consumers uninterrupted services by successfully adopting these ideas.*

**KEYWORDS:** *Functional Isolation, Redundancies, Resilience, Security System.*

---

### INTRODUCTION

Failure domains and redundancies are essential ideas in Site Reliability Engineering (SRE), which is used to develop robust systems. Redundancies relate to having backup systems or duplicate components to maintain continuity in the event of failures, while failure domains refer to discrete regions or components within a system that may fail independently. System availability, downtime, and overall dependability may all be improved by SRE teams by comprehending and using techniques for failure domains and redundancies. An overview of failure domains and redundancies in SRE will be given in this introduction.

First of all, failure domains are necessary to identify and take into consideration probable system failure sites. These failure zones might be physical locations, network segments, hardware parts, or any other element that has the potential to result in sporadic failures. SRE teams may apply techniques to isolate and lessen the effect of failures by identifying and comprehending the failure domains. Implementing fault-tolerant techniques, load balancing across several domains, and planning for resilience and fault isolation are all included in this.

Second, redundancies are essential to maintaining the availability and resilience of the system. Having redundant parts or backup processes that can take over in the case of a breakdown is known as redundancy. This may include using redundant hardware, having many data centers, or duplicating data in several places. SRE teams can lessen the effects of failures and provide users continuous service by implementing redundancy. Additionally, redundancies make it possible to undertake maintenance and updates without affecting how the system functions as a whole[1]–[3].

SRE teams may develop more robust and resilient architectures by designing systems with failure domains and redundancies in mind. Organisations may increase availability, fault tolerance, and user experience by taking into account probable failure sites, putting policies in place to isolate and mitigate failures, and adding redundancy. We will examine several approaches and best practices in this post for determining failure domains, adding redundancies, and developing systems that can endure failures and continue to function. SRE teams may raise the reliability and resilience of their systems by comprehending these ideas and their real-world implementations.

## DISCUSSION

So far, we've discussed how to create systems that can compartmentalize attack fallout and adapt their behavior in reaction to assaults. System design must have redundant components and discrete failure zones to accommodate entire component failures. These strategies should help to mitigate the effects of setbacks and prevent total collapse. Critical component failures must be minimized in particular since any system that relies on them runs the risk of failing completely.

By combining the following strategies, you may develop a balanced solution for your organisation rather than trying to avoid every failure at all times:

1. Separate systems into distinct realms of failure.
2. Reduce the likelihood that a single root cause will have an impact on components across different failure domains.
3. To replace the failing ones, create redundancy resources, parts, or processes.

## Failure Domains

An example of blast radius control is a failure domain. Failure domains accomplish functional isolation by dividing a system into several equivalent but totally independent copies as opposed to architecturally separating by role, place, or time.

**Functional Isolation:** To its customers, a failure domain seems to be a single system. Any of the various partitions may take over for the whole system in the event of a failure, if required. A partition can only support a portion of the system's capacity since it only has a small portion of the system's resources. Operating failure domains and preserving their isolation takes continual work, in contrast to managing role, location, and temporal separations. Failure domains improve system robustness in a manner that conventional blast radius restrictions cannot.

Because not all failure domains are often impacted by a single incident at once, failure domains can shield systems from worldwide damage. A large event, however, has the potential to affect many failure domains or possibly all in severe circumstances. For instance, you may consider the HDDs or SSDs that make up a storage array's foundation as failure domains. The storage system as a whole continues to work even if one device fails because it makes a new data copy somewhere else. Further storage device failures might lead to data loss in the storage system if there aren't enough spare devices to maintain data copies after a significant number of devices fail.

**Data Isolation:** You must be ready for the potential that the data source or certain failure domains may include inaccurate data. In order for each failure domain instance to operate

independently of the other failure domains, each instance requires its own copy of the data. We advise using a two-pronged strategy to establish data separation. You may initially limit how data updates can go into a failure domain. Only when new data passes all validation tests for common and secure modifications is it accepted by a system. A breakglass mechanism<sup>10</sup> may let fresh data to enter the failure domain and certain exceptions are escalated for reason. You are thus more likely to stop hackers or software flaws from making unwanted alterations.

Take ACL modifications, for instance. An empty ACL might be created by a human error or a flaw in the software used to generate the ACLs, which could result in everyone being denied access. A system problem might result from such an ACL update. Similar to this, an attacker may attempt to increase their influence by amending an ACL to include a "permit all" provision. Individual Google services often provide an RPC endpoint for signaling and ingesting new data. The programming frameworks discussed in Chapter 12 contain APIs for versioning data snapshots and determining the accuracy of those snapshots. The reasoning behind the programming framework's determination of whether fresh data is secure may be used by client applications. Quality checks for data updates are implemented through centralized data push services. The data push services determine how to package the data, where to get it from, and when to transmit the finished product. Google uses per-application quotas to rate-limit global updates in order to avoid automation from leading to a broad outage. We forbid operations that alter a number of applications simultaneously or that drastically alter an application's capabilities in a short amount of time.

In addition, systems become more robust to losing access to configuration APIs since they may still utilise the stored configuration when this capability is enabled. In case the most current data is ever damaged, many of Google's systems keep older data around for a short while. Another instance of defence in depth that contributes to long-term resilience is this one.

**Practical aspects:** Even dividing a system into just two failure areas has significant advantages<sup>[4]–[6]</sup>:

1. The ability to do A/B regression is made possible by having two failure domains, which also restricts the blast radius of system modifications to one failure domain. Use one failure domain as a canary and implement a policy that prohibits modifications to both failure domains at the same time to accomplish this feature.
2. Natural catastrophes may be isolated by failure domains that are geographically distant.
3. Different software versions may be used in various failure domains, decreasing the likelihood that a single problem would bring down all servers or destroy all data.

The management of incidents and overall resilience are improved by combining data with functional isolation. By using this strategy, the possibility of unjustified data modifications is reduced. When problems do occur, isolation slows their spread to the various functioning components. This is particularly useful during busy and time-sensitive event response since it provides other defence measures more time to notice and respond. You may independently assess which patches work as planned by simultaneously submitting many candidate fixes to various failure domains. By doing this, you'll be able to prevent mistakenly deploying a hurried update with a flawed "fix" worldwide and further weakening your system as a whole.

Operating expenses are incurred by failure domains. You must save several copies of service settings keyed by failure domain IDs even for a simple service with a few failure domains. To do so, the following is necessary:

1. Consistency of setup is ensured
2. Preventing simultaneous corruption of all setups
3. Hiding the client systems' access to the failure domain separation in order to avoid unintentional coupling to a certain failure domain
4. Partitioning all dependencies could be possible since a modification to a common dependence might unintentionally spread to all failure domains.

It's important to keep in mind that if even one of a failure domain's crucial components fails, the whole domain might collapse. After all, you partitioned the initial system into failure domains in order for the system to continue functioning even if one or more copies of a failure domain entirely fail. Failure domains, however, just move the issue down one level. The danger of total failure of all failure domains is discussed in the section that follows in relation to alternative components.

### Various Component Types

The reliability of a failure domain's components and dependencies taken together is an expression of its resilience. The more failure domains there are, the more resilient the whole system is. However, the operational burden of maintaining more and more failure domains cancels out this enhanced resilience. By delaying or halting the development of new features, you may increase resilience even more while increasing stability. Avoiding additional dependencies also means avoiding any possible failure mechanisms. The frequency of new problems decreases when code updates are discontinued. Even if you stop all feature development, you'll still need to respond to sporadic changes in the state, such as security flaws and rises in customer demand.

Obviously, the majority of organisations can't afford to stop all new feature development. We provide a hierarchy of several strategies for striking a balance between dependability and value in the sections that follow. High capacity, high availability, and minimal reliance are the three main categories of dependability for services.

**High-capacity components:** Your high-capacity service is made up of the elements that you create and maintain as part of your regular business operations. This is so because the primary fleet that serves your consumers is made up of these parts. This is where your service adjusts to sudden increases in customer demand or resource use brought on by new features. The DoS traffic is also absorbed by high-capacity components up to the time when graceful degradation or DoS mitigation take effect. You should concentrate your efforts here initially since these are the parts of your service that are most crucially essential. For instance, you should adhere to the best practises for capacity planning, software and configuration rollouts, and more that are described in Part III of the SRE book and Part II of the SRE workbook.

**High-availability components:** You may reduce these risks by installing duplicates of any high-capacity components in your system that affect all users or have other major wide-reaching effects (like those outlined in the section above). If they give a demonstrably decreased chance of outages, these component copies are said to have high availability.

The copies should be set up with fewer dependencies and a slower pace of change to reduce the likelihood of outages. This strategy lessens the possibility that operational blunders or infrastructural failures may damage the components. You might, for instance, carry out the following:

1. Utilise locally cached data instead of relying on a distant database [7]–[9].
2. To prevent current issues in more recent versions, use older code and configurations.

The operational cost of running high-availability components is low, but it requires extra resources, whose prices scale proportionately with fleet size. A cost-benefit analysis must be done to decide if the high-availability components should support your complete user base or just a piece of it. Set up the graceful degradation features of each high-capacity and high-availability component in the same manner. This enables you to exchange fewer resources for deterioration that is more aggressive.

**Low-dependency components:** A low-dependency service is the next degree of resilience if failures in the high-availability components are intolerable. A different implementation with few dependencies is necessary for low dependence. These low dependencies also have minimum requirements. The total number of tasks, procedures, or services that might fail is as minimal as the demands and expenditures of the organisation will allow. Due to layers of cooperating platforms (virtualization, containerization, scheduling, application frameworks), high-capacity and high-availability services may support enormous user bases and provide complex functionality. While these layers enable services to add or relocate nodes quickly, which aids scalability, they also result in increased outage rates as the cumulative error budgets across the collaborating platforms increase. Low-dependency services, on the other hand, must first simplify their serving stack before they may accept the stack's overall error budget. The serving stack's eventual simplification could necessitate the removal of functionality.

You must assess if it is viable to construct an alternative to a critical component for low-dependency components, provided that the critical and alternative components do not share any failure domains. After all, the likelihood that the same root cause would influence both components is inversely correlated with the success of redundancy. Consider storage as a basic component of a distributed system; in the event that the RPC backends for data storage are down, you may wish to keep local copies of the data. However, it isn't always practicable to save local data copies as a general strategy. The redundant components result in higher operational expenses to sustain them, with little to no added value.

In actuality, you wind up with a small group of low-dependency components that are securely accessible for emergency loads or recovery but have few users, few features, and minimal expenses. Even though the majority of practical features often have several dependencies, a substantially impaired service is preferable than one that is unavailable. Consider a device for which write-only or read-only operations are assumed to be accessible through the network as a small-scale example. These functions in a home security system include writing event logs (write-only) and reading the list of emergency phone numbers (read-only). Disabling the home's internet access is part of an intruder's break-in strategy since it interferes with the security system. You may set up the security system to utilise a local server that uses the same APIs as the distant service in order to prevent this kind of failure. The local server updates the distant service, sends event logs to local storage, and tries again if an attempt is unsuccessful. Requests



to search up emergency phone numbers are likewise answered by the neighbourhood server. The remote service frequently updates the phone number list. The system is operating as intended, producing records and accessing emergency numbers, according to the home security interface. A low-dependency, concealed landline may furthermore provide dialing capabilities as a fallback for a disabled wireless connection [10].

One of the most terrifying sorts of outages, on a business-scale, is a worldwide network failure since it affects both service functioning and the capacity of responders to resolve the outage. Large networks are dynamically controlled and therefore more susceptible to worldwide outages. It takes careful planning to create an alternate network that completely avoids utilising the same connections, switches, routers, routing domains, or SDN software as in the primary network. This design must be focused on a restricted and particular subset of operational parameters and use cases, enabling you to emphasize clarity and readability. Limiting the features and bandwidth that are accessible also naturally results from aiming for modest capital costs for this seldom utilised network. The outcomes are adequate despite the restrictions. Only the most essential functionalities will be supported, and only for a small portion of the typical bandwidth.

**Controlling Redundancies:** Multiple options are set up for each of their dependents in redundant systems. Attackers may be able to take advantage of the differences between the redundant systems—for instance, by pushing the system towards the less secure option—by managing the decision between these possibilities, which is not always simple. Always keep in mind that a resilient design manages to accomplish both security and dependability without compromising one for the other. If anything, attackers who are contemplating wearing down your system may be discouraged when low-dependency alternatives provide superior protection.

**Failover strategies:** In the event of a backend failure, providing a variety of backends, often via load-balancing technologies, increases resilience. For instance, relying on a single RPC backend is impractical. The system will hang whenever that backend has to restart. As long as all backends provide the same feature behaviors, the system often sees redundant backends as interchangeable for simplicity's sake.

A system should depend on a discrete set of interchangeable backends that provide the appropriate reliability behaviors if it requires various reliability behaviors (for the same set of feature behaviors). To decide which set of behaviors to apply and when, the system itself must use logic, maybe via flags. This allows you complete control over the dependability of the system, particularly during recovery. Compare this method to asking the same high-availability backend for low-dependency behavior. You might stop the backend from trying to access its unreachable runtime dependent by using an RPC option. Your system is still one process restart away from collapse if the runtime requirement also affects startup.

It depends on the circumstances when to switch to a component with greater stability. If automated failover is your aim, you should use the techniques described in Controlling Degradation to handle the variations in available capacity. Such a system turns to throttle and load-shedding strategies adjusted for the substitute component after failover. In certain circumstances, you may need to carefully regulate failover or stabilize fluctuations if you want the system to fail back when the failing component recovers.

### Common Pitfalls

Whether running high availability or minimal reliance alternative components, we have seen several frequent difficulties. For instance, you may eventually come to depend on different parts for everyday operations. Any dependent system that starts using the backup systems will probably overflow them during an outage, turning the backup system into an unforeseen source of denial of service. When the alternative components are not regularly utilised, the issue is the reverse, leading to rot and unexpected failures once they are required. Unchecked reliance growing on other services or quantities of necessary computing resources is another danger. Systems often change as user needs shift and developers include new functionality. As dependents and dependencies increase, systems may begin to utilise resources less effectively. When their intended operational restrictions are not regularly checked and confirmed, high-availability copies may lag behind high-capacity fleets or low-dependency services may lose consistency and repeatability.

It is essential that the integrity and security of the system are not jeopardized by failover to substitute components. The appropriate decision will rely on the specifics of your organisation in the following scenarios:

1. You have a high-availability service that, for security reasons (to guard against new defects), is using code that is six weeks old. However, the same service needs a quick security update. Which risk would you rather take: not implementing the repair or the fix possibly damaging the code?
2. By keeping private keys on local storage, it is possible to reduce the startup requirement of a remote key service that retrieves private keys that decrypt data. Does this method provide an unacceptably
3. High danger to those keys, or can the risk be adequately mitigated by increasing the frequency of key rotation?

You decide that you can save up resources by lowering the number of times that regularly changing data (such as ACLs, certificate revocation lists, or user information) is updated. Despite the possibility that doing so would give an attacker more time to modify the data or allow those modifications to survive longer unnoticed, is it still beneficial to release these resources? Finally, you must take care to stop your system from automatically recovering when it shouldn't. It is acceptable for the same resilience measures to automatically unthrottle the system's performance if they automatically throttled it. The drained system may be quarantined due to a security vulnerability, or your team might be managing a cascade failure, therefore if you implemented a manual failover, don't let automation overrule the failover.

## CONCLUSION

To maintain the resilience and availability of systems, failure domains and redundancies are essential components of Site Reliability Engineering (SRE). Organisations may lessen the effects of failures and ensure continuous operation by comprehending and resolving probable failure spots and adopting redundancies. Failure domains assist SRE teams in identifying and isolating possible systemic failure points. Teams may take steps to lessen the effect of failures by identifying discrete sections or components that can fail on their own. Implementing fault-tolerant methods, load balancing across several domains, and designing for resilience and fault isolation are a few examples of how to achieve this. Organisations may proactively prepare for

anticipated failures and lower the risk of significant disruptions by addressing failure domains. On the other hand, redundancies include having backup systems or duplicate parts to guarantee continuation in the case of a breakdown. Organisations may reduce downtime and continue to provide consumers their services by installing redundancy. This may include redundant hardware, many data centers, or data replication across various sites. In addition to helping to handle faults, redundancies enable maintenance or upgrades to be carried out without affecting system availability. In conclusion, a strong SRE approach must have redundancies, fault-tolerant techniques, and an awareness of failure domains. Organisations may create resilient systems that can endure unforeseen occurrences and continue to deliver services to their consumers by proactively addressing probable failures and guaranteeing backup methods.

## REFERENCES

- [1] K. S. Kiangala and Z. Wang, "An Effective Communication Prototype for Time-Critical IIoT Manufacturing Factories Using Zero-Loss Redundancy Protocols, Time-Sensitive Networking, and Edge-Computing in an Industry 4.0 Environment," *Processes*, 2021, doi: 10.3390/pr9112084.
- [2] A. Frigerio, B. Vermeulen, and K. G. W. Goossens, "Automotive Architecture Topologies: Analysis for Safety-Critical Autonomous Vehicle Applications," *IEEE Access*, 2021, doi: 10.1109/ACCESS.2021.3074813.
- [3] T. Simas, R. B. Correia, and L. M. Rocha, "The distance backbone of complex networks," *J. Complex Networks*, 2021, doi: 10.1093/comnet/cnab021.
- [4] L. Myllyaho, M. Raatikainen, T. Männistö, T. Mikkonen, and J. K. Nurminen, "Systematic literature review of validation methods for AI systems," *J. Syst. Softw.*, 2021, doi: 10.1016/j.jss.2021.111050.
- [5] A. Baklouti, N. Nguyen, F. Mhenni, J. Y. Choley, and A. Mlika, "Improved safety analysis integration in a systems engineering approach," *Appl. Sci.*, 2019, doi: 10.3390/app9061246.
- [6] N. Zhu, J. Marais, D. Betaille, and M. Berbineau, "GNSS Position Integrity in Urban Environments: A Review of Literature," *IEEE Trans. Intell. Transp. Syst.*, 2018, doi: 10.1109/TITS.2017.2766768.
- [7] R. A. Kromer, A. Abellán, D. J. Hutchinson, M. Lato, T. Edwards, and M. Jaboyedoff, "A 4D filtering and calibration technique for small-scale point cloud change detection with a terrestrial laser scanner," *Remote Sens.*, 2015, doi: 10.3390/rs71013029.
- [8] M. J. Cannon, A. M. Keller, H. C. Rowberry, C. A. Thurlow, A. Perez-Celis, and M. J. Wirthlin, "Strategies for Removing Common Mode Failures from TMR Designs Deployed on SRAM FPGAs," *IEEE Trans. Nucl. Sci.*, 2019, doi: 10.1109/TNS.2018.2877579.
- [9] P. J. Edmunds, "The hidden dynamics of low coral cover communities," *Hydrobiologia*,

2018, doi: 10.1007/s10750-018-3609-9.

- [10] P. M. Frank, "Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy. A survey and some new results," *Automatica*, 1990, doi: 10.1016/0005-1098(90)90018-D.

## DESIGN FOR RECOVERY

**Ms. Aashi Agarwal\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:aashiagarwal@presidencyuniversity.in

---

### ABSTRACT:

*Designing for recovery is a crucial component of Site Reliability Engineering (SRE) that helps systems recover from failures or interruptions quickly and efficiently. An overview of planning for recovery in SRE is given in this paper. The process of bringing systems back online after a breakdown or event is referred to as recovery. Implementing techniques, procedures, and practices that help systems recover quickly, reduce downtime, and resume regular operation is known as "designing for recovery" (SRE). Organisations may reduce the impact of failures, increase system availability, and improve the overall dependability of their services by prioritizing recovery as a design element. Effective recovery planning not only guarantees company continuation but also boosts client satisfaction and confidence. This paper offers a high-level overview of the major elements and aspects of planning for recovery, which is a crucial facet of SRE.*

**KEYWORDS:** *Fault Tolerance, Effective Monitoring, Redundancy, Resilience.*

---

### INTRODUCTION

Site Reliability Engineering (SRE), which focuses on creating systems and procedures that can rapidly and efficiently recover from failures, disturbances, or events, emphasizes the core idea of designing for recovery. Maintaining the availability, dependability, and performance of vital systems and services requires the capacity to recover quickly.

The following are some of the topics covered in this paper on planning for recovery:

1. Fault tolerance and resilience is essential for recovery to design systems with fault tolerance and resilience in mind. To ensure that failures are limited and do not affect the whole system, this involves putting redundancy, failover methods, and fault isolation into place.
2. Effective monitoring and event detection systems are essential for the early identification of failures or abnormalities. System metrics, logs, and performance indicators may help SRE teams quickly detect problems and start the recovery process.
3. Clear incident response protocols and escalation channels must be established in order to ensure effective recovery. To simplify the response procedure, this entails defining roles and duties, creating communication channels, and putting incident management technologies to use.

4. Automating recovery procedures and developing thorough runbooks may greatly increase the effectiveness of a recovery. Data restoration, infrastructure provisioning, and service restarts are a few examples of automated recovery processes that may speed up recovery and minimise human labor.
5. Regularly testing recovery mechanisms and recreating failure situations makes it easier to spot flaws and confirm the efficacy of recovery solutions. Organisations may proactively identify and fix possible concerns by doing chaotic engineering experiments and disaster recovery simulations.
6. Learning and post-incident analysis:post-incident analysis is crucial for ongoing development. Organisations may improve their ability to recover from disasters and reduce the chance of similar ones happening again in the future by investigating the reasons, figuring out the relevant variables, and putting preventative measures in place.

The process of designing for recovery needs constant interaction between the development, operations, and business teams. It entails taking into account a number of variables, including the system design, fault tolerance techniques, monitoring tools, automation, and incident response procedures. Designing for recovery is a proactive strategy to lessen the effect of these catastrophes in the modern technological world where system failures or mishaps are unavoidable. It entails putting into place procedures, procedures, and processes that allow systems to recover gracefully, reduce downtime, and resume regular operation with little inconvenience to users.The significance of planning for recovery in SRE, its advantages, and crucial factors in the process will all be covered in this introduction.

### **Why Is Recovery-Oriented Design Important?**

Several factors make designing for recuperation essential:

#### **Minimizing Downtime**

Systems that recover quickly spend less time offline or operating below par. By doing this, the effect on users is lessened, possible revenue loss is lessened, and customer happiness is maintained.

#### **Improved Reliability**

Redundancy and recovery techniques make systems more tolerant of errors, which raises their dependability and availability.Assuring Business Continuity: By planning for recovery, vital business operations may be quickly resumed in the event of an incident, minimizing interruptions to existing operations.

#### **Optimizing Mean Time to Recovery (MTTR)**

A key indicator in incident response is MTTR. By streamlining the recovery process and lowering MTTR, designing for recovery makes it possible to resolve incidents more quickly while minimizing the harm they do to users and the company [1]–[3].

### **Key Factors to Consider When Designing for Recovery**

Several important factors should be considered while planning for recovery:

#### **Redundancy and Fault Tolerance**

To guarantee that essential components have backups or other possibilities, use redundancy techniques. This may include installing redundant systems or duplicating data across several sites.

### **Alerting and Monitoring**

Create reliable monitoring and alerting systems to quickly identify and address issues. Proactive monitoring aids in spotting possible problems before they become more serious, enabling speedier healing measures.

### **Runbooks and automation**

Automate the recovery process and write thorough runbooks that detail the specific steps to take after various sorts of accidents. Automation enhances uniformity and minimizes human labour throughout the recovery process.

### **Testing and Drills**

Test the efficiency of recovery mechanisms on a regular basis using exercises and simulations. This enhances trust in the recovery process while identifying possible gaps and improving recovery techniques.

### **Post-Incident Analysis and Learning**

For a complete understanding of the underlying reasons and relevant variables, do a post-incident study. By using preventative measures and improving recovery techniques, accidents may be avoided in the future. Designing for recovery is a proactive and crucial practice that makes sure systems can recover from failures or interruptions promptly and successfully. Building robust systems that recover quickly, preserve business continuity, and provide a pleasant user experience may be accomplished by organisations by taking into account redundancy, monitoring, automation, testing, and learning from events.

## **DISCUSSION**

Failures in contemporary distributed systems may take many different forms and can be caused by malevolent intent as well as inadvertent mistakes. Even the most secure and robust systems need human intervention to be recovered when they are subject to accumulating mistakes, uncommon failure modes, or hostile attacks by attackers.

Unexpected complexity might arise in the process of restoring a broken or compromised system to a stable and secure condition. For instance, reverting to an unstable release can reintroduce security flaws. The deployment of a new version to fix a security flaw can cause reliability problems. These kinds of risky mitigations involve much more delicate decisions. When selecting how rapidly to deploy updates, for instance, a speedy rollout increases your chances of beating attackers but also reduces the amount of testing you can conduct on it. It's possible that you'll release new code with serious stability flaws [4]–[6].

It's not ideal to start worrying about these details and how unprepared your system is to manage them during a tense security or reliability crisis. Only deliberate design choices can provide your system with the dependability and adaptability required to natively handle a range of recovery demands. This chapter discusses a few design guidelines that have helped us set up our systems to support recovery efforts. Numerous of these concepts are applicable at various sizes, ranging

from planet-scale systems to firmware environments inside of specific devices. We'll go through several circumstances that cause a system to need recovery before we get into design ideas to make it easier. These situations may be divided into four groups at their core: random mistakes, unintentional errors, malevolent behavior, and software problems.

### **Random Errors**

Every piece of physical hardware used to construct distributed systems fails. Random mistakes are caused by the unpredictability of physical equipment and the environment in which they function. The possibility of random mistakes occurring in a distributed system rises as the system's physical hardware base expands. Additionally, ageing hardware produces more faults. Some random mistakes are simpler to fix than others. One of the easiest failures to deal with is total failure or isolation of a certain system component, such a vital network router or power supply. Addressing short-lived corruption brought on by unexpected bit flips or long-lived corruption brought on by a failed instruction on one core in a multicore CPU is more difficult. These mistakes are particularly sneaky when they happen in silence.

Modern digital systems may also experience random failures due to external, fundamentally unexpected occurrences. It's possible to abruptly and permanently lose a certain component of the system due to a tornado or earthquake. The provision of electrical power to one or more machines may be jeopardized by a power station or substation failure, a UPS or battery abnormality, or another circumstance. As a result, a voltage sag or swell may be introduced, which may cause memory corruption or other transitory faults.

### **Accidental Errors**

All distributed systems are run, directly or indirectly, by people, and people make errors. Accidental errors are defined as mistakes made by people who had good intentions. The likelihood of human mistake varies depending on the job. Generally speaking, the mistake rate rises as a task's complexity rises. An internal investigation of Google outages from 2015 to 2018 showed that a significant portion (though not all) of failures were brought on by a human action that was taken unilaterally and without first passing through any technical or procedural safety checks. You must take into account how human error could happen across the full stack of tools, systems, and job processes in the system lifecycle since humans might make mistakes in regard to any aspect of your system. Accidental mistakes might also have an unexpected, external influence on your system. For instance, a backhoe being used for unrelated work could accidentally chop through a fiber-optic connection.

### **Errors in software**

The mistake types we've covered so far can be fixed with software updates or design modifications. All faults can be resolved by software, with the exception of software bugs, to paraphrase a well-known adage and its corollary. Mistakes produced during the creation of software are basically simply a unique, delayed example of inadvertent mistakes. You'll need to repair the flaws in your code since it will have them. You may handle issues by using several fundamental design concepts that have been extensively studied, such as testing, code review, and verifying the inputs and outputs of dependent APIs. Software flaws may seem like other kinds of faults. For instance, automation without a safety check may alter production in an abrupt and dramatic way, simulating a malevolent actor. Programme flaws may amplify other sorts of



mistakes, such as sensor problems that provide unexpected numbers that the programme is unable to manage or unusual behavior that seems to be a hostile attack when users work around a broken system.

### **Malicious Actions**

Humans may purposefully try to undermine your systems. These individuals might be privileged insiders who are maliciously informed. Malicious actors are any group of people who are intentionally trying to undermine the security safeguards and dependability of your system(s), or who may be trying to simulate random, unintentional, or other types of failures. Automation may lessen but not completely replace the requirement for human interaction. The size of the organisation responsible for managing your distributed system must grow along with it as its size and complexity increase (preferably, sublinearly). The possibility that one of the people working there may betray the confidence you have in them also increases at the same time.

These trust breaches may be the result of an insider abusing their legitimate control over the system by accessing user data unrelated to their profession, leaking or disclosing trade secrets, or even purposefully attempting to bring about a protracted outage. The individual may sometimes make poor choices, really want to hurt others, become a victim of social engineering, or even come under external pressure. Malicious errors may also be added when a system is compromised by a third party. Regardless of whether the malicious actor is an insider or a third-party attacker who obtains system credentials, mitigating measures for systems are the same[7]–[9].

### **Design Principles for Recovery**

Based on our extensive knowledge of distributed systems, the parts that follow provide some recovery design guidelines. This is not a comprehensive list, but we will provide suggestions for more reading. These guidelines don't simply apply to large corporations like Google, but to a variety of organisations. In general, it's crucial to be open-minded about the breadth and range of issues that might develop while planning for rehabilitation. In other words, concentrate on being prepared to recover from mistakes rather than thinking about how to categorise complex edge situations of failures.

### **Design to Go as Quickly as Possible (Guarded by Policy)**

There is a lot of pressure to restore your system to its intended operating condition as quickly as feasible after a compromise or a system outage. However, the methods you use to modify systems fast run the danger of introducing the incorrect adjustments too soon and aggravating the problem. Additionally, if your systems have been deliberately hacked, hasty recovery or cleaning efforts may result in additional issues, such as alerting an enemy that they have been located.<sup>6</sup> In weighing the tradeoffs inherent in building systems to allow varied rates of recovery, we have discovered a few strategies that work well.

You must be able to alter the state of the system in order to recover your system from any of our four types of errors, or, better yet, to prevent the need for recovery. We recommend designing the update mechanism to operate as quickly as you can imagine it might ever need to operate (or faster, to the extent that it is practical), when building an update mechanism (for example, a software/firmware rollout process, configuration change management procedure, or batch scheduling service). Then, implement controls to limit the pace of change to correspond with

your existing risk and disruption policy. Decoupling your capacity to do rollouts from your rate and frequency rollout regulations has a number of benefits.

Any organization's deployment requirements and procedures evolve over time. For instance, a business may have monthly rollouts in the beginning, never on the weekends or evenings. A change in policy may require challenging refactoring and disruptive code modifications if the rollout system was built to accommodate policy changes. It is far simpler to adapt to unavoidable policy changes that dictate timing and rates of change if the architecture of a rollout system instead clearly isolates the timing and rate of change from the action and substance of that change.

Halfway through a rollout, you could learn something new that changes how you react. Consider that you are distributing an internally designed fix in response to an internally identified security issue. Normally, you wouldn't need to implement this modification quickly enough to run the risk of making your service unstable. If you find out midway through the rollout that the vulnerability is now widely known and being actively exploited in the wild, you may want to speed up the process since your risk assessment may need to alter due to a change in the environment[10].

The risk you're willing to take will eventually alter due to a sudden or unanticipated incident. As a consequence, you should implement a modification as soon as possible. Examples include finding a security flaw or an ongoing compromise. We advise constructing your emergency push system to be nothing more than your ordinary push system on high. Additionally, this implies that your regular rollout system and emergency rollback system are same. We often remark that unproven emergency procedures won't function in an emergency. By making your normal system capable of handling emergencies, you may avoid maintaining two different push systems and ensure that your emergency release system is regularly used.

You'll feel much more confident that your rollout tooling performs as planned if reacting to a difficult scenario just necessitates changing rate limitations in order to swiftly deploy a change. After that, you may concentrate your efforts on other inevitable risks like possible faults in hastily implemented updates or making sure you fix any vulnerabilities that an attacker could have exploited to get access to your systems.

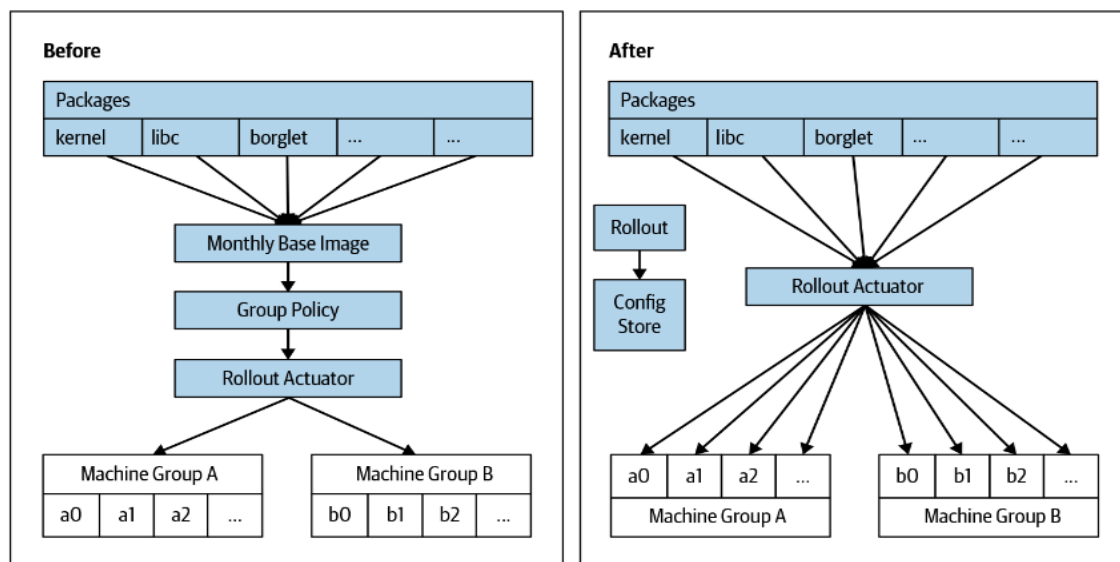
As the implementation of our internal Linux distribution progressed, we discovered these insights. Google used to deploy a "base" or "golden" image with a known set of static files on all the computers in our datacenters. Per computer, there were a few particular changes, including hostname, network settings, and credentials. Every month, as per our protocol, we would deploy a fresh "base" image throughout the fleet. We developed a set of tools and a software update procedure based on that policy and process over a number of years: compile all of the files into a single compressed package, have a senior SRE assess the set of changes, and then progressively update the fleet of computers to the new image.

This policy served as the foundation for our rollout tools, which was created to assign a certain base image to a group of devices. In order to define how to alter that mapping over a period of many weeks, we created the configuration language. Afterward, we utilised a number of ways to add exceptions on top of the underlying picture. Security updates for an increasing number of distinct software packages were one exception; as the list of exclusions became longer, it became less necessary for our tools to follow a monthly cadence.

In response, we made the decision to renounce the presumption that the base image would be updated monthly. Each software package was assigned a more precise release unit that we created. On top of the current rollout technique, we also created a brand-new API that defined the precise collection of packages to install on each individual computer. Figure 1 illustrates how this API detached the programme that specified various aspects:

- a) The rollout and the pace of change for each package were both intended
- b) The configuration database that listed every machine's current configuration
- c) The rollout actuator controls how each machine is updated.

We were able to separately design each component as a consequence. We then constructed a rollout system to monitor and maintain the individual rollouts of each package, and we repurposed an existing config database to define the configuration of all the packages deployed to each machine. A considerably larger variety of release speeds for various packages might be enabled by divorcing the image construction from the monthly rollout process. Additionally, certain test machines might run the most recent releases of all the software while still maintaining a reliable and consistent rollout to the majority of the fleet's workstations. Even better, separating the policy from the system allowed for new systemic applications.



**Figure 1: The evolution of our workflow for deploying packages to machines.**

We currently use it to routinely disseminate to the whole fleet a portion of thoroughly screened files. By simply tweaking certain rate limitations and authorizing the release of one sort of package to go more quickly than usual, we can also utilise our standard tools for emergency releases. The ultimate product was more straightforward, practical, and secure.

### Limit Your Dependencies on External Notions of Time

Time, or the regular time of day as shown by tools like wall clocks and wristwatches, is a kind of condition. Any site where your system includes wall-clock time might possibly compromise your capacity to finish a recovery since you're often unable to change how your system perceives the passage of time. Unexpected system behaviors might result from discrepancies between the

moment you start your recovery attempt and the last time the system was working properly. For instance, unless you design the recovery process to take the original transaction date into account when verifying certificates, a recovery that requires replaying digitally signed transactions may fail if certain transactions are signed by expired certificates. If your system's sense of time relies on an external concept that you don't manage, it may be significantly more likely to cause security or reliability problems. This pattern appears in a variety of problems, including software mistakes like Y2K, Unix epoch rollovers, and unintentional mistakes when developers set certificate expiry dates that are "not their problem anymore." If a network attacker is able to take over, clear-text or unauthenticated NTP connections can pose a danger. A code smell that suggests you could be making a time bomb is the presence of a fixed date or time offset.

## CONCLUSION

Different facets of creating systems for recovery were covered in this chapter. We provided an explanation of why systems should be flexible in terms of how quickly they roll out changes. This flexibility enables you to roll out changes slowly when it is possible to do so and prevent coordinated failures, as well as quickly and confidently when you need to accept greater risk in order to achieve security goals. Building dependable systems requires the ability to roll back changes, but sometimes you may need to stop rollback to versions that are unreliable or sufficiently old. To reliably restore the system to any previously functional state and make sure that its current state complies with your security requirements, it is essential to comprehend, monitor, and reproduce the state of your system to the greatest extent possible through software versions, memory, wall-clock time, and other factors. Emergency access, as a last choice, enables responders to stay connected, evaluate a system, and lessen the situation. The path to recoverable systems is paved by thoughtfully managing policy vs procedure, the central source of truth against local functions, and the anticipated state versus the actual state of the system. This also fosters resilience and strong daily operations.

## REFERENCES

- [1] A. De Lucia, V. Deufemia, C. Gravino, and M. Risi, "Design pattern recovery through visual language parsing and source code analysis," *J. Syst. Softw.*, 2009, doi: 10.1016/j.jss.2009.02.012.
- [2] G. Antoniol, G. Casazza, M. Di Penta, and R. Fiutem, "Object-oriented design patterns recovery," *J. Syst. Softw.*, 2001, doi: 10.1016/S0164-1212(01)00061-9.
- [3] G. Rasool and D. Streitfdert, "A Survey on Design Pattern Recovery Techniques," *J. Comput. Sci. Issues*, 2011.
- [4] N. Krishnamoorthy *et al.*, "Engineering principles and process designs for phosphorus recovery as struvite: A comprehensive review," *J. Environ. Chem. Eng.*, 2021, doi: 10.1016/j.jece.2021.105579.
- [5] G. Rasool, I. Philippow, and P. Mäder, "Design pattern recovery based on annotations," *Adv. Eng. Softw.*, 2010, doi: 10.1016/j.advengsoft.2009.10.014.
- [6] B. Huang, K. Pu, P. Wu, D. Wu, and J. Leng, "Design, selection and application of energy recovery device in seawater desalination: a review," *Energies*. 2020. doi: 10.3390/en13164150.

- [7] M. M. Smith, J. D. Aber, and R. Rynk, "Heat Recovery from Composting: A Comprehensive Review of System Design, Recovery Rate, and Utilization," *Compost Science and Utilization*. 2017. doi: 10.1080/1065657X.2016.1233082.
- [8] R. J. Urbanic and W. H. El Maraghy, "Using axiomatic design with the design recovery framework to provide a platform for subsequent design modifications," *CIRP J. Manuf. Sci. Technol.*, 2009, doi: 10.1016/j.cirpj.2008.09.019.
- [9] A. Sadowska, F. Świdorski, and W. Laskowski, "Osmolality of components and their application in the design of functional recovery drinks," *Appl. Sci.*, 2020, doi: 10.3390/app10217663.
- [10] G. Rasool, P. Maeder, and I. Philippow, "Evaluation of design pattern recovery tools," in *Procedia Computer Science*, 2011. doi: 10.1016/j.procs.2010.12.134.

---

## MITIGATING DENIAL-OF-SERVICE ATTACKS

Dr. Ganpathi Chandankeri\*

\*Associate Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:chandankeri@presidencyuniversity.in

---

### ABSTRACT:

*To maintain the availability and stability of systems and services, Site Reliability Engineering (SRE) must include the mitigation of denial-of-service (DoS) assaults. DoS attacks try to disable or disrupt systems by flooding them with unapproved traffic or by taking advantage of holes in their defences. An overview of methods for preventing DoS attacks in an SRE setting is given in this abstract. Putting preventative measures, detection tools, and reaction plans into place may help mitigate DoS attacks by reducing their effects.*

**KEYWORDS:** *Dos, Network Traffic, Traffic Analysis, Rate Throttling.*

---

### INTRODUCTION

A crucial component of site reliability engineering (SRE), which ensures the availability and dependability of systems, is the mitigation of denial-of-service (DoS) assaults. DoS attacks try to overtax or deplete system resources, disrupting service for authorised users. Organisations can lessen the effects of DoS attacks and preserve continuous operation by employing preventative measures and efficient mitigation solutions. Here are some essential measures to lessen DoS attacks:

**Network-Level Protections:** Put in place network-level security measures including rate limiters, firewalls, and intrusion detection systems (IDS). These technologies may be used to restrict the amount of requests coming from a single source, detect abnormal traffic patterns, and block them in order to shield the target system from malicious activity.

**Content Delivery Networks (CDNs):** Make use of CDNs to distribute traffic and provide caching. By dispersing the traffic over many servers and absorbing and filtering it, CDNs lighten the stress on the origin system. Additionally, they provide defence against different DoS assaults, including volumetric assaults[1]–[3].

**Load Balancers:** To uniformly distribute incoming traffic across many servers or instances, use load balancers. Load balancers aid in load distribution and guard against DoS attacks that overload a single server. Additionally, they offer dynamic scaling, which enables the addition of extra resources to cope with growing demand.

**Application-Level Protections:** Use user authentication techniques, rate restriction, and CAPTCHA as application-level security measures. By restricting the amount of requests made by any individual user or IP address, rate limiting helps keep the system from being overloaded with too much traffic. By distinguishing between human and artificial traffic, CAPTCHA

challenges may lessen the effect of bot-driven DoS assaults. User authentication assists in identifying trustworthy users and thwarts harmful efforts.

**Traffic Analysis and Anomaly Detection:** To track and spot odd traffic patterns, use systems for anomaly detection and traffic analysis. By seeing unexpected increases in traffic or unusual behavior, these technologies may assist in real-time detection and mitigation of DoS assaults.

**Cloud-Based Services:** Utilise the DoS prevention services provided by cloud providers that are cloud-based. These services often come with built-in DoS mitigation tools including rate limitation, traffic filtering, and automatic attack detection.

**Incident Response Planning:** Create an incident response strategy that is especially suited to addressing DoS assaults. This strategy should contain predetermined actions for isolating compromised systems, turning on mitigations, and engaging the appropriate parties in the response procedure. To make sure the strategy is effective, test and update it often.

**DDoS Mitigation Services:** Consider collaborating with a specialized DDoS mitigation service provider for bigger organisations or those that are more at danger. These service providers provide cutting-edge traffic filtering and mitigation methods, using their knowledge and infrastructure to assist defend against significant DoS assaults.

**Regular Monitoring and Testing:** Keep an eye out for any indications of a DoS assault by continuously monitoring system logs, network traffic, and performance indicators. To proactively discover and resolve possible problems, perform vulnerability assessments and penetration testing on a regular basis.

**Collaboration and Information Sharing:** Keep up with the most recent DoS attack patterns, methodologies, and mitigation tactics by cooperating and exchanging information with other organisations and security groups. Sharing information and expertise may improve how well organisations as a whole defend against DoS assaults.

Organisations may greatly lessen the effects of DoS assaults and guarantee the availability of their systems by putting these mitigation methods into place. It is crucial to use a multi-layered strategy that combines preventative measures at the network, application, and infrastructure levels with proactive monitoring and testing to remain watchful. To combat this always changing threat environment, DoS assaults must be mitigated using a proactive and all-encompassing strategy.

### **Understanding DoS attack types**

SRE teams can create effective defences against DoS assaults by knowing the many kinds, including volumetric, application layer, and distributed attacks.

### **Network traffic analysis**

Utilising tools and methods for network traffic analysis aids in spotting unusual traffic patterns that might point to a possible DoS attack. This involves keeping track on traffic volume, bandwidth use, and request pattern irregularities.

### **Load balancing and scaling**

Traffic may be distributed properly and the effects of DoS attacks can be reduced by using load balancing methods and scaling systems either horizontally or vertically. The system can manage

more traffic by distributing the demand among many resources, which lowers the likelihood that it will become overloaded.

### **Rate throttling and rate limiting**

The quantity of requests or connections from certain sources may be managed by using rate restriction and throttling techniques. This provides equitable resource distribution and prevents a single source from overwhelming the system.

### **Content delivery networks (CDNs)**

Utilising CDNs may assist in regionally distributing traffic and mitigating the effects of volumetric assaults. The burden on the origin servers is lessened by CDNs' delivery and caching services that are closer to the end users.

### **Intrusion detection and prevention systems (IDPS)**

Real-time DoS attack detection and prevention are made possible by using IDPS tools and methodologies. Network traffic is monitored by IDPS systems, which spot suspicious patterns or behaviors and block or reduce assaults as necessary.

### **Incident response and recovery**

Effective DoS attack mitigation requires the establishment of comprehensive incident response strategies and processes. To quickly restore services, this entails establishing roles and responsibilities, determining escalation channels, and putting recovery procedures into practice. Organisations may improve their defence against DoS assaults and guarantee the availability and dependability of their systems and services by using these tactics and procedures.

## **DISCUSSION**

Security professionals often consider assault and defence in relation to the systems they guard. However, economics provides more useful terminology for a typical denial-of-service attack: the adversary tries to raise the demand for a certain service over the capacity of that service's supply. As a consequence, the service no longer has the capacity to accommodate its genuine consumers. The company will then have to choose between suffering downtime (and related financial losses) until the assault ends or incurring even more costs by trying to absorb the attack. While some sectors are more commonly the target of DoS attacks than others, any service might be subject to one of these assaults. DoS extortion, a kind of financial assault where the adversary threatens to interrupt services until money is paid, often targets random targets.

### **Strategies for Attack and Defense**

In order to accomplish their objectives, attackers and defenders must make effective use of their limited resources. It is vital to begin developing a defensive plan by comprehending your opponent's tactics so you may identify vulnerabilities in your defences before they can. With this knowledge, you may build defences against well-known assaults and create systems that are flexible enough to swiftly mitigate brand-new threats.

### **Attacker's Strategy**

An attacker must concentrate on making the most of their little resources in order to outperform their target. A cunning foe may be able to sabotage the operations of a stronger foe.



Most services have a number of dependents. Think about how a typical user request might proceed:

1. The IP address of the server that must accept user traffic is revealed via a DNS query.
2. The request is sent across the network to the service frontends.
3. The user request is interpreted by the service frontends.
4. Database capability is offered by service backends for personalised replies.

The service will be interrupted by an assault that is successful in stopping any of those processes. The majority of inexperienced attackers will try to deliver a deluge of network traffic or application requests. A highly skilled attacker can create requests that are more expensive to process, for instance by exploiting the search feature that is available on many websites.

A determined adversary will create techniques for harnessing the power of several computers in what is known as a distributed denial-of-service (DDoS) assault since a single machine is seldom sufficient to interrupt a major service (which is often supported by many machines). An attacker may conduct an amplification assault or infect susceptible computers and use them as part of a botnet to carry out a DDoS attack[4]–[6].

### **Defender's Approach**

A well-equipped defence may deflect assaults by simply oversupplying their whole stack, but doing so comes at a high cost. Power-hungry devices occupy costly datacenters, and it is impossible to supply always-on capacity to withstand the most powerful assaults. While services established on a cloud platform with substantial capacity may have the option of automated scaling, defenders often need to use other practical, affordable methods to safeguard their services.

Engineering time must be taken into consideration while choosing your finest DoSdefence approach; you should give priority to those that will have the most effect. While it may be appealing to concentrate on fixing the problem from yesterday, recency bias may lead to priorities that change very quickly. Instead, we advise using a threat model strategy to focus your efforts on the dependence chain's weakest link. Threats may be compared based on how many computers an attacker would need to take over in order to interrupt users' daily activities.

### **Making Defence Designs**


An ideal attack concentrates all of its resources on a single resource that is limited, such network bandwidth, the CPU or memory of an application server, or a backend service like a database. Protecting each of these resources as effectively as you can should be your aim.

Attack traffic gets more concentrated and more costly to counteract as it travels farther into the system. Layered defences, where each layer defends the one below it, are hence a crucial design element. Here, we look at the design decisions that result in systems that are defendable at two fundamental design layers: common infrastructure and individual services.

### **Defendable Architecture**

Most services use similar infrastructure, such as network load balancers, application load balancers, peering capacity, and load balancers for networks. Providing common defences in the

shared infrastructure makes perfect sense. High bandwidth assaults may be slowed down by edge routers, defending the backbone network. Attacks that flood networks with packets may be throttled by network load balancers to safeguard application load balancers. Before the traffic reaches service frontends, application load balancers may throttle attack-specific assaults.



**Figure 1: A DDoS attack on a site protected by Project Shield, as seen from (top) the perspective of the individual site, and (bottom) the perspective of the Project Shield load balancers.**

Since you only need to capacity-plan inner layers for DoS attack types that can get past the defences of outer levels, layering defences is often more cost-effective. Attack traffic should be stopped as soon as feasible to save bandwidth and computing power. For instance, you may block suspect traffic before it has a chance to use up internal network capacity by installing ACLs at the network edge. Similar cost reductions may be obtained by deploying caching proxies close to the network edge, which also reduces latency for authorized users.

Putting defences in place for shared infrastructure also offers a beneficial economies of scale. Shared defences enable you to cover a variety of services with a single provisioning, even while it may not be cost-effective to provide strong defence capabilities for each particular service. For instance, Figure 1 demonstrates how an assault on one site resulted in traffic levels that were much higher than usual for that site but remained manageable when compared to the traffic levels seen by all of the websites covered by Project Shield. Similar bundling strategies are used by commercial DoS mitigation services to provide customers a practical answer.

Similar to how a magnifying glass may capture the force of the sun to start a fire, a DDoS assault that is extremely severe might overrun a datacenter's capabilities. Any defence plan must make sure that a scattered attack's force cannot be concentrated on a single component. This kind of overload may be avoided by using network and application load balancers to continuously monitor incoming traffic and distribute it to the closest datacenter with available capacity.

Using anycast, a method in which an IP address is broadcast from many places, you may protect shared infrastructure without depending on a reactive system. By using this strategy, each venue

draws visitors from surrounding users. As a consequence, a distributed assault will be spread out over many places across the globe and won't be able to concentrate its force on a single datacenter.

### **Defendable Services**

The way a website or application is designed may have a big influence on how well a service is protected. Although the greatest defence is to make sure that the service degrades gracefully in overload situations, there are a number of simple improvements that may be done to increase resistance to assault and provide substantial cost savings during regular operation:

#### **Make use of caching proxies**

Repeated requests for content may be handled by proxies with the use of the Cache-Control and associated headers, eliminating the need for each request to go through the application backend. The majority of static photos go under here, and maybe the main page as well[7]–[9].

#### **Avoid unnecessary application requests**

It's ideal to reduce the amount of queries required since each one uses server resources. It is more effective to serve all of the little icons on a web page as one (bigger) image, a process known as spriting, if the page includes many small icons. As an added bonus, fewer requests from actual users will result in fewer false positives for harmful bot detection.

#### **Reduce the egress bandwidth.**

Traditional assaults aim to fill up the ingress bandwidth, however it is also feasible for an attack to fill up your bandwidth by making a lot of requests for resources. Users' website load times will be sped up by resizing pictures to only take up the space they really need. Another method is to rate restrict or deprioritize unavoidably huge answers.

#### **Mitigating Attacks**

Even while a defendable architecture makes it possible to survive many DoS assaults, you could still require active defences to stop more significant or complex attacks.

**Monitoring and Alerting:**Mean time to detection (MTTD) and mean time to repair (MTTR) account for the majority of the time it takes to resolve an outage. The server's CPU use may increase as a result of a DoS attack, or the programme may run out of memory while queuing requests. You must track the request rate in addition to CPU and memory utilisation in order to quickly identify the main problem.

A clear indicator of an assault may be provided to the incident response team by alerting on abnormally high request rates. Make sure your pager notifications can be taken action on, however. It is often preferable to just absorb an assault if it is not harming users. Only when demand surpasses service capacity and automatic DoSdefences have kicked in do we advise notifying.The rule of only warning when taking human action could be necessary also applies to network-layer threats. Many synflood attempts may be neutralized, however if syncookies are found, it could be necessary to issue a warning.<sup>8</sup> Similar to this, high-bandwidth assaults only merit a page if a connection gets overloaded[10].

**Graceful Degradation:**If absorbing an assault isn't an option, you should try to minimise the user-facing effect. Network ACLs may be used to restrict suspicious traffic during a significant

---

attack, acting as an efficient switch to instantly reduce attack volume. It's crucial to sometimes let real traffic that fits the attack signature through while avoiding completely blocking suspect traffic in order to maintain system visibility. Throttles may not be enough since a cunning attacker might fake actual traffic. Additionally, you may provide priority to important traffic by using quality-of-service (QoS) controls. If necessary, bandwidth may be released to higher QoS queues by using a lower QoS for less crucial traffic, such as batch copies.

## CONCLUSION

Even if they don't think they'll be the target of a DoS assault, every online service should be prepared for them. Each organisation has a maximum on the amount of traffic it can handle, and the defender's job is to efficiently neutralize assaults that exceed this capacity. It's crucial to keep in mind the financial limitations of your DoS defences. The cheapest course of action is seldom to simply absorb an assault. Use cost-effective mitigation strategies instead, beginning with the design stage. When under assault, weigh all of your alternatives, including banning a troublesome hosting provider which could also involve a few actual users or going offline briefly and informing your users of the issue. Also keep in mind that the "attack" could not have been planned. Collaboration between different teams is necessary to implement defences at each tier of the serving stack. DoS defence may not be a high concern for all teams. Focus on the financial savings and organizational simplifications a DoS mitigation system can provide to win their support. Instead of requiring to withstand the most significant assaults at every stage of the stack, capacity planning may concentrate on the actual user demand. Using a web application firewall (WAF), the security team may concentrate on new threats by filtering known dangerous requests. The same approach may stop efforts at exploitation if you find application-level vulnerabilities, giving the development team time to create a fix.

## REFERENCES

- [1] A. Ezenwe, E. Furey, and K. Curran, "Mitigating denial of service attacks with load balancing," *J. Robot. Control*, 2020, doi: 10.18196/jrc.1427.
- [2] R. Singh, S. Tanwar, and T. P. Sharma, "Utilization of blockchain for mitigating the distributed denial of service attacks," *Secur. Priv.*, 2020, doi: 10.1002/spy2.96.
- [3] J. Mölsä, "Mitigating denial of service attacks: A tutorial," *J. Comput. Secur.*, 2005, doi: 10.3233/JCS-2005-13601.
- [4] N. Schweitzer, A. Stulman, A. Shabtai, and R. D. Margalit, "Mitigating Denial of Service Attacks in OLSR Protocol Using Fictitious Nodes," *IEEE Trans. Mob. Comput.*, 2016, doi: 10.1109/TMC.2015.2409877.
- [5] M. Usman, M. Qaraqe, M. R. Asghar, and I. Shafique Ansari, "Mitigating distributed denial of service attacks in satellite networks," *Trans. Emerg. Telecommun. Technol.*, 2020, doi: 10.1002/ett.3936.
- [6] D. Alves, "A Strategy for Mitigating Denial of Service Attacks on Nodes with Delegate Account of Lisk Blockchain," in *ACM International Conference Proceeding Series*, 2020. doi: 10.1145/3390566.3391684.
- [7] G. Rajakumaran and N. Venkataraman, "Moving target defense strategy for mitigating denial of service attack in the public cloud environment," *Int. J. Eng. Adv. Technol.*, 2019,

doi: 10.35940/ijeat.F8077.088619.

- [8] O. S. Akanji, O. A. Abisoye, and M. A. Iliyasu, "Mitigating Slow Hypertext Transfer Protocol Distributed Denial of Service Attacks in Software Defined Networks," *J. Inf. Commun. Technol.*, 2021, doi: 10.32890/JICT2021.20.3.1.
- [9] E. Bertino, M. Kantarcioglu, C. G. Akcora, S. Samtani, S. Mittal, and M. Gupta, "AI for Security and Security for AI," in *CODASPY 2021 - Proceedings of the 11th ACM Conference on Data and Application Security and Privacy*, 2021. doi: 10.1145/3422337.3450357.
- [10] S. Kumar and R. Amin, "Mitigating distributed denial of service attack: Blockchain and software-defined networking based approach, network model with future research challenges," *Secur. Priv.*, 2021, doi: 10.1002/spy2.163.

## CASE STUDY: DESIGNING, IMPLEMENTING, AND MAINTAINING A PUBLICLY TRUSTED CA

**Mr. Narayana Gopalakrishnan\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id: gopalakrishnan@presidencyuniversity.in

### ABSTRACT:

*In the field of cybersecurity, creating, implementing, and maintaining a publicly trustworthy Certificate Authority (CA) is a challenging and crucial undertaking. In order to create trust and ensure safe communication in the digital world, digital certificates must be issued and managed by a CA. An overview of the factors to take into account and difficulties that may arise while creating, putting into practice, and maintaining a publicly trustworthy CA are given in this abstract. The infrastructure, rules, and practices required to guarantee the reliability, security, and availability of certificate issuing and management operations must be defined as part of the architecture of a publicly trusted CA. This entails creating safe methods for storing private keys, putting in place reliable authentication procedures, and setting up effective auditing and monitoring tools. Adherence to industry standards and best practices is necessary for the implementation of a publicly trustworthy CA. Gaining the confidence of users and dependent parties requires adherence to security measures like the Certificate Policy (CP) and the Certificate Practice Statement (CPS). In addition to interacting with other PKI ecosystem participants, implementation also entails creating secure channels of communication for certificate issue and revocation.*

**KEYWORDS:** Auditing, Digital Certificates, Monitoring, Operating Systems.

### INTRODUCTION

Continuous work is required to guarantee the reliability and security of the CA infrastructure in order to maintain a publically trustworthy CA. This entails patching and upgrading software on a regular basis, performing penetration tests and security audits, and swiftly revocation of certificates in the event that breaches or vulnerabilities are found. It is possible to spot any irregularities or certificate abuse with regular monitoring and analysis of certificate issue and revocation actions. To guarantee that its digital certificates are widely recognized and accepted, a publicly trustworthy CA must also build and maintain strong relationships with dependent parties, such as web browsers and operating systems. This entails adhering to the standards and regulations established by certification authorities (CAs) consortiums and root programs.

Security, compliance, and industry standards must be carefully taken into account while establishing, deploying, and maintaining a publicly trusted CA. A publicly trusted CA can provide a safe and reliable basis for digital communications by adhering to best practices and keeping up with new risks, allowing users to comfortably participate in secure online

transactions and communication. A challenging and vital undertaking is creating, establishing, and sustaining a publically trusted Certification Authority (CA). Digital certificates that verify the integrity and validity of websites, digital identities, and other online entities are issued by a publicly trusted CA. A few crucial factors must be taken into account in order to build and retain trust:

### **Infrastructure for security**

Create a strong security system that protects the private keys that are used to sign certificates. To safeguard the private keys and provide strict access restrictions to thwart unauthorized access, use Hardware Security Modules (HSMs).

### **Auditing and Compliance**

Ensure adherence to legal and industry standards, including the Web Trust programme. Regularly submit to independent audits to confirm compliance and show stakeholders that they can be trusted.

### **Management of the Certificate Lifecycle**

Implement a clear procedure for managing the lifespan of certificates, which should include their issue, renewal, revocation, and key rotation. Keep accurate and current records of all issued certificates and the information related to them[1]–[3].

### **Mechanisms for revoking certificates**

Create procedures for quickly revoking compromised or mishandled certificates. For real-time certificate validation, keep CRLs up to date or utilise OCSP (Online Certificate Status Protocol).

### **Hierarchy of Public Key Infrastructure (PKI)**

Create a PKI hierarchy that is effective and scalable to control certificate issuance and validation. To spread the burden and improve security, take into account the deployment of subordinate CAs for specialized purposes or organizational divisions.

### **Security Techniques**

Implement robust security procedures throughout the whole CA infrastructure, including encrypted network connections, safe data storage, and secure backup procedures. To guard against vulnerabilities, update and patch all software and systems on a regular basis.

### **Certificate Openness**

To improve transparency and identify incorrectly issued or fraudulent certificates, take into consideration installing Certificate Transparency (CT) logs. All certificates issued by the CA are listed in the CT logs, enabling monitoring and verification.

### **Observance of Browser and Industry Standards**

To preserve confidence in the CA, make sure compliance with browser and industry regulations. This entails adhering to the Baseline Requirements established by the CA/Browser Forum and staying current with best practices and standards as they change.

### **Recovery from disasters and business continuity**

Create thorough disaster recovery and business continuity strategies to guarantee that, in the case of a catastrophic failure or natural catastrophe, the CA's activities can be rapidly resumed. Test these strategies often to confirm their efficacy.

### **Continuous Monitoring and Improvement**

Implement ongoing monitoring and improvement methods to find and fix performance problems, security vulnerabilities, and new threats. Keep abreast of market changes and advancements to guarantee that the CA's infrastructure stays reliable and secure.

A complete grasp of security procedures, industry standards, and legal requirements is necessary for designing, putting into practice, and sustaining a publicly trustworthy CA. It requires constant dedication to compliance, security, and continuing improvement. Organisations may create and maintain a CA that is respected by users, operating systems, and browsers by adhering to five crucial requirements, guaranteeing the safe and dependable functioning of digital certificates.

### **DISCUSSION**

By providing certificates for Transport Layer Security (TLS), S/MIME, and other typical distributed trust situations, publicly trustworthy certificate authority serve as trust anchors for the internet's transport layer. They comprise the group of CAs that devices, operating systems, and browsers automatically trust. As a result, there are many security and dependability issues to take into account while creating and maintaining a publicly trusted CA.

A CA must meet a variety of standards that cover various platforms and use cases in order to gain public confidence and keep it. Publicly trusted CAs must, at the very least, submit to audits against industry standards like WebTrust and those established by groups like the European Telecommunications Standards Institute (ETSI). Additionally, publicly trusted CAs must comply with the CA/Browser Forum Baseline Requirements. A typical publicly trusted CA spends at least one-fourth of each year on these audits, which evaluate logical and physical security controls, processes, and practices. Additionally, in order for a CA to be trusted by default, it must satisfy the specific criteria of the majority of browsers and operating systems. CAs must be flexible and open to changing infrastructure or processes as needs evolve.

Most businesses depend on other parties to get public TLS certificates, code signing certificates, and other sorts of certificates that call for widespread user confidence, therefore it's unlikely that your company will ever need to create a publicly trusted CA. This case study's objective is to illustrate some of our results that may be relevant to projects in your environment rather than to demonstrate how to create a publically trustworthy CA. Important conclusions included the following:

1. The environment was made more secure by our choice of programming language and our decision to handle data produced by other parties via segmentation or containers.
2. For resolving core dependability and security concerns, rigorous testing and hardening of code both code we produced ourselves and third-party code was essential.
3. We simplified the design and automated manual processes to make our infrastructure safer and more dependable.



4. We were able to develop validation and recovery procedures that help us better prepare for a catastrophe in advance thanks to our understanding of our threat model.

### What Made a Publicly Trusted CA Necessary?

Over time, our company's requirements for a publicly reputable CA altered. We acquired all of our public certificates from a third-party CA in the early years of Google. We sought to address three issues that were inherent in this method:

- a. **Reliance on third parties:** We need rigorous validation and control over certificate issuance and management since some business requirements call for a high degree of confidence, such as providing clients with cloud services. We were unclear of whether outside parties could maintain a high quality of safety, even if we conducted required audits inside the CA ecosystem. Our opinions on security have been reinforced by notable security breaches at publicly reputable CAs.
- b. **Need for automation:** Global users may access Google's hundreds of company-owned domains. We intended to secure every domain we hold and often rotate certificates as part of our ubiquitous TLS efforts (see Example: Increasing HTTPS use). Additionally, we wanted to make it simple for consumers to get TLS certificates. It was challenging to automate the procurement of fresh certificates since many external publicly trusted CAs had extendable APIs or offered SLAs that were insufficient for our purposes. As a consequence, several manual procedures that are prone to mistake were used in the certificate request process.
- c. **Cost:** Cost study revealed it would be more cost-effective to develop, deploy, and manage our own CA rather than continuing to get certificates from third-party root CAs given the millions of TLS certificates Google wished to employ for its own web sites and on behalf of customers[4]–[6].

### The Build or Buy Decision

Once Google made the decision to run a publicly trusted CA, we had to choose between purchasing commercial software and developing our own software to run the CA. In the end, we chose to build the CA's core ourselves, with the flexibility of incorporating open source and proprietary solutions as needed. There were other determining variables, however the following were the main drivers of this choice:

- a. **Transparency and validation:** Commercial CA solutions often lacked the supply chain or degree of auditability for code that we need for such crucial infrastructure. Even though it utilised some proprietary code from third parties and was integrated with open source libraries, designing and testing our own CA software boosted our trust in the system we were constructing.
- b. **Integration capabilities:** We integrated with Google's secure essential infrastructure to make it easier to build and maintain the CA. For instance, using only one line in a configuration file, we might configure Spanner to do routine backups.
- c. **Flexibility:** New efforts were being developed by the larger internet community to strengthen ecosystem security. Two classic instances are domain validation utilising DNS, HTTP, and other protocols, and certificate transparency, which allows for the monitoring and

auditing of certificates. We wanted to be early adopters of these efforts, and the fastest way to provide this flexibility was via a bespoke CA.

### **Considerations for Design, Implementation, and Maintenance**

We developed a three-layer tiered architecture for our CA's security, with each layer handling a separate aspect of the certificate issuing procedure: certificate request processing, Registration Authority operations (routing and logic), and certificate signing. Microservices with clear roles make up each tier. A dual trust zone design that handles untrusted input in a distinct setting from important operations was also developed by our team. This division establishes well defined limits that encourage comprehension and reviewability. The design also makes mounting an attack more difficult since components have restricted functionality, which limits the functionality an attacker who obtains access to a component may change. The attacker would need to go through more audit points in order to obtain access to more information.

Simplicity is a critical design and implementation guideline for any microservice. We progressively rework each component with simplicity in mind during the course of the CA. We rigorously test and validate all of our code internal and third-party as well as our data. When doing so would increase safety, we containerize code as well. In further depth, this section explains our strategy for tackling security and dependability via sound design and implementation decisions.

### **Choosing a Programming Language**

The design included a crucial decision on the programming language to be used for system components that receive untrusted input at random. In the end, we elected to create the CA in a mixture of Go and C++, and we choose the appropriate language for each subcomponent depending on its function. Both Go and C++ display good performance, are compatible with tried-and-true cryptographic libraries, and offer a robust ecosystem of frameworks and tools to carry out routine tasks.

Go provides some extra benefits for security if the CA handles arbitrary input since it is memory-safe. Examples of untrusted input entering the CA include Certificate Signing Requests (CSRs). CSRs might originate from a user of the internet (perhaps even a hostile actor), or from one of our internal systems, which may be reasonably secure. We wanted to adopt a memory-safe language that offered additional protection since there is a lengthy history of memory-related flaws in code that decodes DER (Distinguished Encoding Rules, the encoding standard used for certificates)<sup>6</sup>. Go was appropriate.

Although C++ is not memory-safe, it has high system compatibility, particularly for several of Google's key infrastructure components. We execute this code in a safe zone and verify every data before it enters that zone in order to make it secure. For instance, when processing CSRs, we parse the request in Go first, then pass it on to the C++ subsystem for the same action, comparing the outcomes thereafter. Processing is stopped if there is a discrepancy.

### **Complexity vs Comprehensibility**

As a precaution, we consciously decided to design our CA with fewer functionalities than the entire range of choices provided by the standards (see *Designing Understandable Systems*). We issued certificates for typical web services with widely used characteristics and extensions as our

main use case. Our analysis of available commercial and open source CA software choices revealed that these solutions' efforts to account for esoteric qualities and extensions that we didn't require caused the system to become complicated, making the programme more difficult to evaluate and prone to mistakes. As a result, we chose to create a CA that was simpler to comprehend and had fewer features so that we could more readily audit the intended inputs and outputs.

We are always working to make the CA architecture simpler so that it is easier to comprehend and manage. In one instance, we discovered that our design had produced an excessive number of distinct microservices, leading to higher maintenance expenses. Although we desired the advantages of a modular service with well-defined limits, we discovered that consolidating certain system components was easier. In a different instance, we discovered that our ACL checks for RPC calls were carried out manually in every instance, opening the door for developer and reviewer mistake. To centralize ACL checks and remove the chance of new RPCs being introduced without ACLs, we refactored the programme[7]–[9].

### **Securing Open Source and Third-Party Components**

Our bespoke CA depends on third-party code in the form of for-profit modules and open-source libraries. This code needs to be verified, strengthened, and containerized. As a starting point, we concentrated on the several popular and well-known open source applications that the CA employs. Even open source products that are widely utilised in security-related situations and that come from sources with solid security histories are prone to flaws. Each received a thorough security evaluation from us, and we contributed patches to fix any concerns we discovered. As far as feasible, we also put all open source and third-party components through the testing process described in the next section. We also have some additional layering protection against defects or malicious code insertions thanks to the deployment of two secure zones one for processing sensitive activities and one for handling untrusted data. The aforementioned CSR parser utilizes free X.509 libraries and operates as a microservice in a Borg container in the untrusted zone. This adds an additional layer of defence against problems with this code.

We also needed to protect closed-source, proprietary code from other parties. Using a hardware security module (HSM), a specialized cryptographic processor, offered by a commercial vendor to serve as a vault safeguarding the CA's keys, is necessary to run a publicly trusted CA. For the vendor-provided code that communicates with the HSM, we needed to add another layer of validation. The types of tests we could carry out were constrained, as they are with many vendor-supplied systems. We did the following to safeguard the system against issues like memory leaks:

1. Because we were aware that the inputs or outputs may be dangerous, we developed the components of the CA that required to communicate with the HSM libraries cautiously.
2. The third-party code was executed using nsjail, a simple process isolation tool.
3. We informed the seller of any difficulties we discovered.

### **Testing**

We create unit and integration tests to account for a variety of situations in order to maintain project cleanliness. As part of the development process, team members are required to create

these tests, and peer reviews make sure that this standard is followed. We test for negative situations in addition to predicted behavior. We produce test certificate issuance circumstances every few minutes, some of which satisfy excellent requirements and others of which have serious flaws. We specifically verify, for instance, that correct error messages trigger alarms when an unauthorized user issues a document. We can immediately undertake high-confidence end-to-end testing on all new CA software installations since we have a library of both positive and negative test situations.

We also benefit from the advantages of integrated automated code testing on both pre-submit and post-build artefacts by using Google's centralized software development toolchains. Tricorder, our static analysis platform, reviews all code changes at Google, as was covered in Integration of Static Analysis in the Developer Workflow. In order to find frequent problems, we additionally expose the CA's code to a number of sanitizers, including AddressSanitizer (ASAN) and ThreadSanitizer (see Dynamic Programme Analysis). In addition, we specifically fuzz the CA code (see Fuzz Testing).

### Validation of Data

The worst faults a CA may make, apart from losing vital materials, are issuance errors. We attempted to construct our systems so that human judgement cannot affect validation or issuance. As a result, we can concentrate on the accuracy and reliability of the CA code and infrastructure. A system's expected behavior is confirmed by continuous validation (see Continuous Validation). At several phases of the issuing process, we automatically run certificates through linters to put this idea into practice in Google's publicly trusted CA. The linters look for patterns of errors, such as confirming that `subject:commonName` has a proper length or that certificates have a valid lifespan. After the certificate has been verified, we add it to the Certificate Transparency records so that the public may continue to verify it. We additionally use several separate logging systems as a last line of defence against fraudulent issuing, which we can reconcile by contrasting the two systems entry by entry to assure consistency. Before entering the log repository, these logs are signed for further security and potential subsequent validation[10].

### CONCLUSION

An example of infrastructure with stringent security and dependability criteria is certificate authorities. When infrastructure is implemented according to the recommended practices stated in this book, security and dependability may improve over time. Although you should include these concepts into a design from the beginning, you should also utilise them to enhance systems as they become older.

### REFERENCES

- [1] S. Khan, Z. Zhang, L. Zhu, M. Li, Q. G. K. Safi, and X. Chen, "Accountable and Transparent TLS Certificate Management: An Alternate Public-Key Infrastructure with Verifiable Trusted Parties," *Secur. Commun. Networks*, 2018, doi: 10.1155/2018/8527010.
- [2] B. Li, J. Lin, Q. Wang, Z. Wang, and J. Jing, "Locally-Centralized Certificate Validation and its Application in Desktop Virtualization Systems," *IEEE Trans. Inf. Forensics Secur.*, 2021, doi: 10.1109/TIFS.2020.3035265.

- [3] A. Singla and E. Bertino, "Blockchain-Based PKI solutions for IoT," in *Proceedings - 4th IEEE International Conference on Collaboration and Internet Computing, CIC 2018*, 2018. doi: 10.1109/CIC.2018.00-45.
- [4] B. Li, F. Li, Z. Ma, and Q. Wu, "Exploring the security of certificate transparency in the wild," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2020. doi: 10.1007/978-3-030-61638-0\_25.
- [5] J. Q. Lin, J. W. Jing, Q. L. Zhang, and Z. Wang, "Recent advances in PKI technologies," *Journal of Cryptologic Research*. 2015. doi: 10.13868/j.cnki.jcr.000095.
- [6] M. D. Campos, A. Camacho, K. Pereda, K. Santana, I. Calix, and T. W. Fong, "Attitudes towards gambling, gambling problems, and treatment among hispanics in imperial county, CA," *J. Gambl. Stud.*, 2016, doi: 10.1007/s10899-015-9585-3.
- [7] J. H. Koo, B. H. Kim, and D. H. Lee, "Authenticated public key distribution scheme without trusted third party," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2005. doi: 10.1007/11596042\_95.
- [8] D. Gisolfi, "Self-sovereign identity: Why blockchain?," *Blockchain Pulse: IBM Blockchain Blog*, 2018.
- [9] V. M. Shah and V. V Kapadia, "A Review on Modern Methods of Encryption: Tendencies and Challenges," *Int. J. Adv. Res. Comput. Sci. Softw. Eng.*, 2014.
- [10] G. Aguilar, L. Lundsberg, S. Baer, C. Kennedy, N. Stanwood, and A. Gariepy, "ORAL ABSTRACTS," *Contraception*, 2021, doi: 10.1016/j.contraception.2021.07.018.

## A STUDY TO EVALUATING AND BUILDING FRAMEWORKS

**Dr. Jagdish Godihal\***

\*Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:drjagdishgodihal@presidencyuniversity.in

---

### ABSTRACT:

*This abstract will provide a general overview of the framework evaluation and construction process and will highlight important factors to take into account while making choices. According to a number of criteria, including functionality, performance, scalability, community support, documentation, and compatibility with current systems, the assessment process comprises determining the viability of existing frameworks. It's critical to completely investigate and contrast several frameworks in order to comprehend their advantages, disadvantages, and compatibility with project needs. Building atop a framework starts after it has been decided upon. To do this, the framework must be altered and expanded to accommodate the demands of the project. Developing new modules, incorporating third-party libraries, or changing already existing components are all examples of customization. During the construction phase, it is crucial to take maintainability, compatibility, and the effect on future updates into account.*

**KEYWORDS:** *Compatibility, Dependability, Frameworks, Security.*

---

### INTRODUCTION

Framework evaluation and construction are crucial components of software development and system design. Frameworks provide an organized approach and reusable parts that speed up the development process, better the quality of the code, and improve system performance as a whole. But choosing and creating the best structure needs thorough assessment and consideration of numerous criteria. In this introduction, we'll discuss the value of assessing and creating frameworks as well as the crucial factors to take into account while doing so. Frameworks provide pre-defined structures, libraries, and functions that speed up development and give standardised methods. They act as a basis for creating software applications. They might be either general-purpose frameworks or ones that focus on a particular industry or class of applications. The effectiveness, scalability, and maintainability of the final programme are directly impacted by the evaluation and selection of an appropriate framework.

The project needs, scalability, performance, usability, community support, and compatibility with current systems are crucial factors to take into account while choosing frameworks. The optimal framework for a project is one that supports developer productivity, fits with the organization's technology stack, and is identified after a comprehensive examination [1]–[3]. Organisations may also need to create new frameworks to satisfy certain business objectives in addition to analysing current frameworks. Creating reusable components, coding standards, architectural patterns, and best practices documentation are all necessary steps in the framework-

building process. Custom frameworks provide specialized solutions, guaranteeing consistency across projects, encouraging code reuse, and enhancing the effectiveness of development.

Technical know-how, project requirements comprehension, and sector expertise are all necessary for evaluating and developing frameworks. Finding the best course of action entails doing research, analysis, and prototyping. Organisations may optimize software development, raise the quality of the code, and ensure long-term maintainability by devoting time and effort to this process. We shall examine the numerous factors involved in assessing and creating frameworks in this essay. We'll look at several evaluation standards for already-existing frameworks, talk about how to create unique frameworks, and provide advice on best practices and problems to watch out for. Organisations may make well-informed choices that result in effective software development and effective system designs by realizing the value of reviewing and developing frameworks.

Important factors for assessing and developing frameworks include:

**Project prerequisites:** Make that the chosen framework adheres to the project requirements, which should be precisely defined. Take into account elements like the programming language, database support, needed user interface, and performance objectives.

**Community and Support:** As it directly affects the accessibility of resources, documentation, and community-driven assistance, consider the size and activity of the framework's community. Timely updates, bug patches, and a robust ecosystem of plugins and extensions are all made possible by a thriving community.

**Scalability and Performance:** Evaluate the performance and scalability of the framework. Think about how well it can manage a growing user base, large amounts of data, and system integration. Case studies and performance benchmarks may provide insightful information.

**Flexibility and Extensibility:** Identify the framework's degree of flexibility and extension. Customization and expansion are possible with a solid foundation without sacrificing essential functionality. Examine the availability of plugins, modules, and libraries that can expand the capabilities of the framework.

**Documentation and Learning Curve:** Examine the documentation for the framework's quality and thoroughness. The learning curve for developers is shortened by clear and current documentation, which also encourages best practices and makes maintenance and problem-solving easier.

**Long-Term Maintenance and Support:** Take into account the framework's durability and continuous support. Examine the product's release schedule, version stability, and development team response to bug fixes and security upgrades. Long-term dependability and security are ensured by a well-managed framework.

**Integrity and Compatibility:** Examine the framework's interoperability with current databases, libraries, and systems. Find out if integration will be simple and whether there will be any difficulties along the way.

Organisations may make sure that their software projects are based on a strong basis by carefully reviewing and developing frameworks. A framework that facilitates effective development, achieves project objectives, and permits future expansion is produced by making educated

judgements based on project needs, community support, scalability, adaptability, documentation, and compatibility.

## **DISCUSSION**

Because security and dependability are difficult to retrofit into software, it's critical to take them into consideration from the beginning of the design process. It is difficult and less productive to add these features after a launch, and you may need to alter other core beliefs about the codebase.

Educating developers is the first and most crucial step in lowering security and reliability concerns. Even highly skilled engineers may make errors; for example, security professionals can produce vulnerable code, while SREs can overlook reliability concerns. It's challenging to keep all of the factors and trade-offs involved in creating safe and dependable systems in mind at once, particularly if you're also in charge of writing software.

You may ask SREs and security experts to assess code and software designs rather of relying entirely on developers to check it for security and dependability. This strategy is also not ideal; human code inspections won't uncover every flaw, and no reviewer will detect every security flaw that a prospective attacker may exploit. Reviewers may sometimes be swayed by their own interests or experiences. For instance, they might have a natural tendency to look for novel attack classes, complex design flaws, or intriguing cryptographic protocol flaws; in contrast, reviewing a large number of HTML templates for cross-site scripting (XSS) flaws or examining the error-handling logic for each RPC in an application might be regarded as less exciting.

While not always successful, code reviews do offer other advantages. Developers are encouraged to write their code in a manner that makes the security and reliability characteristics simple to evaluate by having a robust review culture. The tactics for incorporating automation into the development process and for making these qualities clear to reviewers are covered in this chapter. These tactics may free up a team's time to concentrate on other problems and help establish a culture of security and dependability.

The process of evaluating and creating frameworks is crucial to software engineering and development. Frameworks provide reusable parts, libraries, and architectural patterns as a base for creating applications. But choosing and creating the ideal framework for a particular project might be difficult.

### **Frameworks for Securing and Reliable Systems**

Domain-specific invariants are essential to an application's security and dependability. A programme is safe against SQL injection attacks, for instance, if all of its database queries are written entirely by the developer and external inputs are provided through query parameter bindings. If all user input that is entered into HTML forms is correctly escaped or sanitized to eliminate any executable code, a web application may fend against XSS assaults. Theoretically, you can write application code that carefully preserves these invariants to provide safe and trustworthy software. This method becomes almost difficult as the number of needed attributes and the complexity of the codebase increase. It is impossible to expect any developer to be an authority on any of these topics or to write or review code with continual attention to detail.



If every modification must be manually reviewed by humans, those individuals will struggle to preserve global invariants as reviewers often lose sight of the global context. A reviewer must be acquainted with all transitive callers of a function if they are to know which function parameters are handed user input by callers and which arguments solely carry developer-controlled, reliable values. It's doubtful that reviewers will be able to maintain this condition over time. Dealing with security and dependability in popular frameworks, languages, and libraries is a preferable strategy. Libraries should ideally just offer an interface that prevents the creation of code with common types of security flaws. Each library or framework may be used by a variety of applications. This engineering method scales better when domain experts address a problem by removing it from all the apps the framework supports. Using a centralized hardened framework decreases the possibility of new vulnerabilities surfacing in comparison to manual evaluation. Of course, no framework can completely guard against security flaws; attackers may still unearth previously undiscovered attack types or errors in the framework's implementation. But rather than patching the codebase as a whole, you may address a new vulnerability in one or more specific places.

To provide an exact illustration: On both the OWASP and SANS lists of common security vulnerabilities, SQL injection (SQLI) takes the top rank. In our experience, these kinds of vulnerabilities become unimportant when you utilise a robust data library like TrustedSqlString (see SQL Injection Vulnerabilities: TrustedSqlString). These presumptions are made clear by types, which the compiler automatically enforces[4]–[6].

### **Advantages of Framework Use**

For security (authentication and authorization, logging, data encryption), and dependability (rate limiting, load balancing, retry logic), the majority of apps share comparable basic pieces. It is costly and results in a patchwork of various issues in each service to develop and maintain such building blocks from scratch for each service.

By allowing developers to customize a single building block, rather than taking into consideration all of the security and reliability factors influencing a particular capability or feature, frameworks promote code reuse. For instance, a developer doesn't need to worry about the veracity of the information they supply since the framework will determine which credentials from the incoming request are crucial for permission. A developer may also define which data should be recorded without thinking about replication or storage. Frameworks also make updating simpler since you just need to apply changes once.

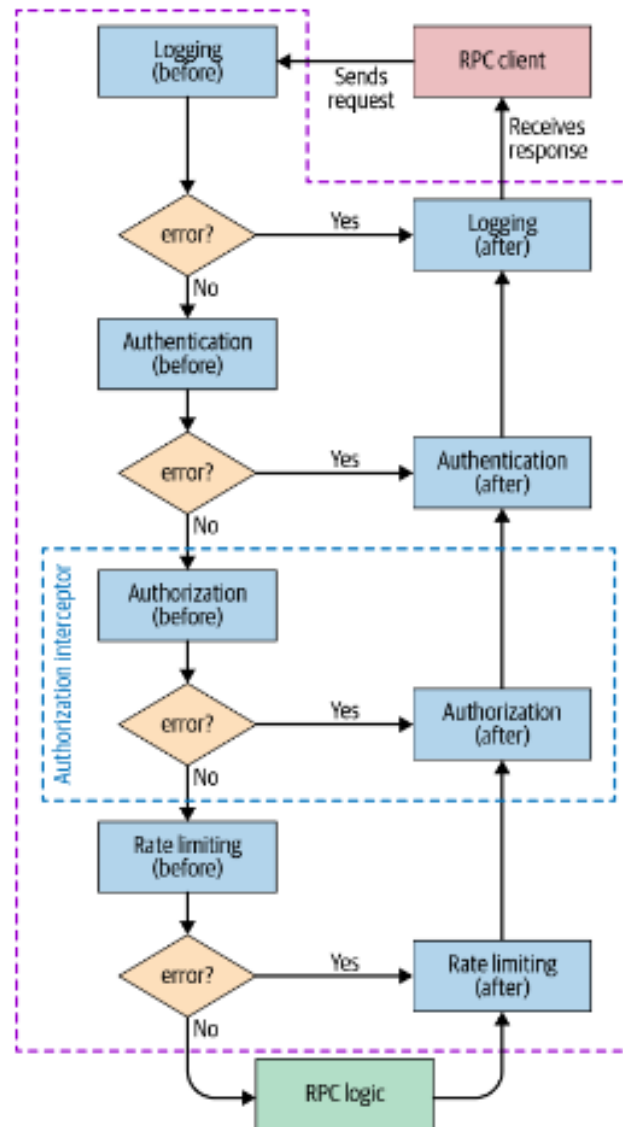
Building a culture of security and dependability may benefit from using frameworks since it increases productivity for all developers within an organisation. The framework's building blocks should be designed and developed by a group of domain specialists rather than by each team separately, which would be significantly less effective. For instance, if the security team is in charge of cryptography, the other teams may use their expertise. Developers utilising the frameworks may concentrate on the business logic of an application without having to worry about any of the frameworks' internal features.

Frameworks provide tools that are simple to integrate with, which further boosts efficiency. Examples of fundamental operational metrics that frameworks may automatically output are the overall number of requests, the number of unsuccessful requests broken down by error type, or the processing stage delay. This information may be used to create automated dashboards for

service monitoring and alerts. Frameworks also make it simpler to integrate with load-balancing infrastructure, enabling a service to automatically reroute traffic away from stressed-out instances or launch new service instances as necessary. Services that are developed on top of frameworks thus have much greater dependability.

By distinctly separating business logic from common functions, frameworks also make it simple to think about the code. This gives developers greater confidence when asserting a service's security or dependability. In general, frameworks result in less complexity since it's simpler to adhere to industry best practices when the code across several services is more standard. It's not always a good idea to create your own frameworks. Reusing current solutions is often the best course of action. An established and well-liked framework like Tink (described in Example: Secure cryptographic APIs and the Tink crypto framework) is something you may utilise instead of building and implementing your own cryptographic framework, as practically every security expert will tell you to avoid doing.

It's crucial to assess a framework's security posture before choosing to use it. In order to include the most recent security changes for any code on which your code relies, we also advise adopting actively maintained frameworks and consistently upgrading your code dependencies. A potential framework design based on specified interceptors that are in charge of each of the aforementioned phases is shown in Figure 1. Interceptors could also be used for unique steps. Each interceptor specifies a step to be taken both before and after the RPC logic actually runs. Each stage has the ability to report an error situation that stops the execution of further interceptors. The after stages of each interceptor that has already been called, however, are carried out in reverse order when this happens. The framework in between the interceptors may transparently carry out additional tasks, such as exporting performance data or error rates. The logic carried out at each level is clearly separated by this design, which increases simplicity and dependability.



**Figure 1: A control flow in a potential framework for RPC backends: the typical steps are encapsulated in predefined interceptors and authorization is highlighted as an example.**

In this example, the *before* stage of the logging interceptor could log the call, and the *after* stage could log the status of the operation. Now, if the request is unauthorized, the RPC logic doesn't execute, but the "permission denied" error is properly logged. Afterward, the system calls the authentication and logging interceptors' *after* stages (even if they are empty), and only then does it send the error to the client. Interceptors share state through a *context object* that they pass to each other. For example, the authentication interceptor's *before* stage can handle all the cryptographic operations associated with certificate handling (note the increased security from reusing a specialized crypto library rather than reimplementing one yourself). The system then wraps the extracted and validated information about the caller in a convenience object, which it adds to the context. Subsequent interceptors can easily access this object [7], [8].

The framework can then use the context object to track request execution time. If at any stage it becomes obvious that the request won't complete before the deadline, the system can automatically cancel the request. You can increase service reliability by notifying the client quickly, which also conserves resources. A good framework should also enable you to work with dependencies of the RPC backend for example, another backend that's responsible for storing logs. You might register these as either soft or hard dependencies, and the framework can constantly monitor their availability. When it detects the unavailability of a hard dependency, the framework can stop the service, report itself as unavailable, and automatically redirect traffic to other instances.

Sooner or later, overload, network issues, or some other issue will result in a dependency being unavailable. In many cases, it would be reasonable to retry the request, but implement retries carefully in order to avoid a *cascading failure* (akin to falling dominoes).<sup>1</sup> The most common solution is to retry with an *exponential backoff*.<sup>2</sup> A good framework should provide support for such logic, rather than requiring the developer to implement the logic for every RPC call. A framework that gracefully handles unavailable dependencies and redirects traffic to avoid overloading the service or its dependencies naturally improves the reliability of both the service itself and the entire ecosystem. These improvements require minimal involvement from developers<sup>[9], [10]</sup>.

## CONCLUSION

To guarantee the efficacy and efficiency of systems and activities, Site Reliability Engineering (SRE) evaluates and builds frameworks as a critical activity. Organisations may construct reliable and scalable systems that match their unique demands by carefully evaluating existing frameworks and taking reasoned choices throughout the building process. A comprehensive review of the alternatives is required when evaluating frameworks, taking functionality, compatibility, performance, security, and simplicity of use into account. Through this assessment process, organisations may choose frameworks that meet their needs and provide the capabilities required to support their operations. Building frameworks entails creating and putting into action special frameworks catered to the particular requirements of the organisation. This method requires a thorough understanding of the framework's required functionality, scalability, and extensibility. Additionally, it entails taking future growth and changes into account by taking maintainability, modularity, and flexibility into account.

## REFERENCES

- [1] G. D. Sepich-Poore, L. Zitvogel, R. Straussman, J. Hasty, J. A. Wargo, and R. Knight, "The microbiome and human cancer," *Science*. 2021. doi: 10.1126/science.abc4552.
- [2] T. O. Olawumi and D. W. M. Chan, "Green-building information modelling (Green-BIM) assessment framework for evaluating sustainability performance of building projects: a case of Nigeria," *Archit. Eng. Des. Manag.*, 2021, doi: 10.1080/17452007.2020.1852910.
- [3] A. Al-Sakkaf, T. Zayed, and A. Bagchi, "A Sustainability Based Framework for Evaluating the Heritage Buildings," *Int. J. Energy Optim. Eng.*, 2020, doi: 10.4018/ijeoe.2020040105.
- [4] M. Abdel-basset, A. Gamal, R. K. Chakraborty, M. Ryan, and N. El-saber, "A comprehensive framework for evaluating sustainable green building indicators under an

- uncertain environment,” *Sustain.*, 2021, doi: 10.3390/su13116243.
- [5] S. BuHamdan, A. Alwisy, B. Barkokebas, A. Bouferguene, and M. Al-Hussein, “A multi-criteria lifecycle assessment framework for evaluating building systems design,” *J. Build. Eng.*, 2019, doi: 10.1016/j.jobe.2019.02.010.
- [6] M. Kong, M. Lee, H. Kang, and T. Hong, “Development of a framework for evaluating the contents and usability of the building life cycle assessment tool,” *Renew. Sustain. Energy Rev.*, 2021, doi: 10.1016/j.rser.2021.111475.
- [7] R. Gupta, M. Gregg, and R. Cherian, “Developing a new framework to bring consistency and flexibility in evaluating actual building performance,” *Int. J. Build. Pathol. Adapt.*, 2020, doi: 10.1108/IJBPA-04-2019-0032.
- [8] K. Sperling, F. Hvelplund, and B. V. Mathiesen, “Evaluation of wind power planning in Denmark - Towards an integrated perspective,” *Energy*, 2010, doi: 10.1016/j.energy.2010.06.039.
- [9] P. Fazio, H. S. He, A. Hammad, and M. Horvat, “IFC-Based Framework for Evaluating Total Performance of Building Envelopes,” *J. Archit. Eng.*, 2007, doi: 10.1061/(asce)1076-0431(2007)13:1(44).
- [10] G. Tang, Z. Zhao, J. Yu, Z. Sun, and X. Li, “Simulation-based framework for evaluating the evacuation performance of the passenger terminal building in a Ro-Pax terminal,” *Autom. Constr.*, 2021, doi: 10.1016/j.autcon.2020.103445.

## A BRIEF INTRODUCTION ON SECURITY AND RELIABILITY

**Dr. Mohammad Shahid Gulgundi\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:mohammadshahid@presidencyuniversity.in

---

### ABSTRACT:

*Any digital infrastructure or system must have both security and dependability. The security of sensitive data, preserving service availability, and protecting against hostile attacks are crucial in today's technologically advanced and linked society. The paper explores the difficulties that organisations have in establishing security and dependability, including the dynamic threat environment, intricate technology ecosystems, and the necessity to strike a compromise between usability and practicality and security measures. It emphasizes the need of putting into place an extensive and proactive strategy to deal with these difficulties, including risk assessments, security audits, personnel training, continuing monitoring, and testing. It also covers the advantages of placing a high priority on security and dependability, including safeguarding private data, preserving client confidence, abiding by legal obligations, and preventing loss of money and goodwill.*

**KEYWORDS:** *Data Breaches, Dependability, Reliability, Unauthorized Access.*

---

### INTRODUCTION

Security and dependability are two essential SRE pillars that guarantee the stability, accessibility, and protection of systems and data. Organisations nowadays must put a high priority on security measures, develop dependable systems, and retain customer trust in order to protect against cyber-attacks, preserve consumer confidence, and avoid service interruptions. The vital significance of security and dependability in the context of digital systems is explored in this abstract. It identifies the primary obstacles that organisations have in obtaining and sustaining these qualities and offers alternative solutions, including best practices. Security refers to a broad set of practices and policies used to safeguard systems, networks, and data from unauthorized access, breaches, and hostile actions. It entails putting into place reliable authentication and access control systems, encryption methods, intrusion detection and prevention systems, and thorough security policies. Organisations may reduce the risk of data breaches, theft, and unauthorized access by taking a proactive and all-encompassing approach to security. This helps to guarantee the privacy, integrity, and accessibility of their systems and data.

The continuous and uninterrupted operation of systems and services is the emphasis of reliability, on the other hand. It entails creating architectures that are resilient to errors, putting fault-tolerant and redundant systems in place, and making sure that backups are made on schedule and disaster recovery plans are in place. In order to identify problems quickly, minimise downtime, and guarantee high availability for users, reliability also includes monitoring and warning systems. The interconnectedness of security and dependability is also highlighted in the

abstract. Since attacks, breaches, and interruptions are less likely to occur, a secure system is often more dependable. Similar to this, a secure system has a higher probability of being dependable since it is better equipped to deal with errors and reduce risks. Security and dependability are essential building blocks of every digital system. In order to protect their operations, data, and systems, organisations must give priority to these factors. Organisations may increase their resilience, defend against threats, and provide their customers a safe and dependable experience by taking a proactive and all-encompassing strategy, putting in place strong security controls, and guaranteeing continuous availability. An overview of the significance of security and dependability in SRE will be given in this introduction. Implementing safeguards to guard against unauthorized access, data breaches, and malicious activity is a key component of SRE security. It includes a number of elements, including vulnerability management, incident response, encryption, authentication, and authorization. SRE teams can stop data breaches, reduce risks, and guarantee regulatory compliance by giving security first priority.

On the other side, reliability focuses on creating and maintaining systems that provide consumers continuous, uninterrupted services. It entails creating resilient structures, putting redundancy into place, and setting up systems to deal with failures, outages, and interruptions. The goals of reliability measures are to reduce downtime, respond to events swiftly, and provide a smooth user experience [1]–[3]. In SRE, security and dependability go hand in hand. Systems are susceptible to assaults and breaches that might jeopardize critical data, harm their reputation, and interrupt services if they lack security. However, in the absence of dependability, even secure systems may often go down or have performance difficulties, frustrating users and causing them to lose faith in the system. Organisations may benefit from SRE's integration of security and dependability in a number of ways.

### **Security of Sensitive Information**

Strong security measures are put in place to guarantee the privacy, availability, and integrity of sensitive data, shielding the organisation and its users from security breaches and unauthorized access.

### **Reducing online threats**

Before they have a chance to do any harm, possible cyber threats may be identified and mitigated with the use of proactive security techniques like vulnerability scanning, threat intelligence, and intrusion detection systems.

### **Increasing Customer Trust**

Users' faith in the organization's security and dependability procedures increases their trust in the organization's services. For retaining long-term client connections and luring in new ones, this trust is necessary.

### **Adherence to regulations**

Compliance with industry rules like the General Data Protection Regulation (GDPR) or the Payment Card Industry Data Security Standard (PCI DSS) is ensured through the implementation of security controls and dependable systems. Compliance aids in avoiding both possible legal repercussions and financial fines.

## Reduction of interruptions and downtime

Proactive monitoring and robust systems with built-in redundancy reduce downtime and service interruptions. As a result, the user experience is enhanced, customers are happier, and sales are up. The core components of site reliability engineering are security and dependability. Organisations may safeguard their systems and data, reduce risks, foster customer trust, and provide dependable services by giving these factors top priority. For any organisation to succeed and endure in the constantly changing digital environment, security and dependability principles must be woven into the foundation of SRE.

## DISCUSSION

### Common Security Vulnerabilities

Even with continual attempts to train developers and provide code review, the bulk of security vulnerabilities in big codebases are concentrated in a small number of classes. Lists of typical vulnerability classes are published by OWASP and SANS. The list [#top\\_onezero\\_most\\_common\\_vulnerability\\_r](#) provides some possible framework-level mitigation strategies for each of the OWASP's top 10 vulnerability risks.

OWASP top 10 vulnerability	Framework hardening measures
[SQL] Injection	<a href="#">TrustedSQLString</a> (see the following section).
Broken authentication	Require authentication using a well-tested mechanism like OAuth before routing a request to the application. (See <a href="#">Example: Framework for RPC Backends</a> .)
Sensitive data exposure	Use distinct types (instead of strings) to store and handle sensitive data like credit card numbers. This approach can restrict serialization to prevent leaks and enforce appropriate encryption. Frameworks can additionally enforce transparent in-transit protection, like HTTPS with LetsEncrypt. Cryptographic APIs such as <a href="#">Tink</a> can encourage appropriate secret storage, such as loading keys from a cloud key management system instead of a configuration file.
XML external entities (XXE)	Use an XML parser without XXE enabled; ensure this risky feature is disabled in libraries that support it. <sup>2</sup>
Broken access control	This is a tricky problem, because it's often application-specific. Use a framework that requires every request handler or RPC to have well-defined access control restrictions. If possible, pass end-user credentials to the backend, and enforce an access control policy in the backend.
Security misconfiguration	Use a technology stack that provides secure configurations by default and restricts or doesn't allow risky configuration options. For example, use a web framework that does not print error information in production. Use a single flag to enable all debug features, and set up your deployment and monitoring infrastructure to ensure this flag is not enabled for public users. The <code>environment</code> flag in Rails is one example of this approach.
Cross-site scripting (XSS)	Use an XSS-hardened template system (see <a href="#">Preventing XSS: SafeHtml</a> ).
Insecure deserialization	Use deserialization libraries that are built for handling untrusted inputs, such as <a href="#">Protocol Buffers</a> .
Using components with known vulnerabilities	Choose libraries that are popular and actively maintained. Do not pick components that have a history of unfixed or slowly fixed security issues. Also see <a href="#">Lessons for Evaluating and Building Frameworks</a> .
Insufficient logging & monitoring	Instead of relying on ad hoc logging, log and monitor requests and other events as appropriate in a low-level library. See the logging interceptor described in the previous section for an example.

**Figure 1: Top 10 most common vulnerability risks according to OWASP.**

The type-based safety strategy we discussed in the part before is not only for preventing SQL injection. To lessen the likelihood of cross-site scripting vulnerabilities in web applications, Google used a more complicated variation of the same concept. XSS flaws fundamentally arise when a web application presents suspect input without using the proper sanitization. For instance, an application may interpolate a value for `$address` that is under the control of an



attacker into an HTML fragment like `div>${address}/div>`, which is shown to a different user. The attacker may then set `address` to execute arbitrary code in the context of another user's page by setting it to the script `"exfiltrate_user_data();"`.

The equivalent of binding query parameters does not exist in HTML. Instead, before being included into an HTML page, suspect values must be suitably sanitized or escaped. Application developers must also interpret values differently based on the context in which they occur since various HTML attributes and elements have distinct meanings. For instance, the `javascript:` technique might result in code execution from an attacker-controlled URL. Distinct types for values intended for distinct contexts, for as `SafeHtml` to represent the contents of an HTML element and `SafeUrl` to represent URLs that are safe to go to, may be included in a type system to capture these needs. The constructors provided for each type uphold the contracts; each type is a (immutable) wrapper over a string. The trusted codebase, which is in charge of guaranteeing the application's security attributes, is made up of the constructors[4]–[6].

For various use situations, Google has developed a variety of building libraries. Builder methods that need the right type for each attribute value and `SafeHtml` for the element contents may be used to create individual HTML elements. The `SafeHtml` contract for more complex HTML is guaranteed by the template system with strong contextual escaping. It does the following:

#### **Parses the partial HTML in the template**

Determines each replacement point's context.e.g., appropriately escaping or sanitising untrusted string values, or requiring the programme to give in a value of the right type.If you have the following Closure Template, for instance:

```
{template .foo kind="html"}<script src="{ $url}"></script>{/template }
```

trying to use a string value for `$url` will fail:

```
templateRendered.setMapData(ImmutableMap.of("url", some_variable));
```

Instead, the developer has to provide a `TrustedResourceUrl` value, e.g.:

```
templateRenderer.setMapData(
  ImmutableMap.of("x", TrustedResourceUrl.fromConstant("/script.js")) ).render();
```

You shouldn't include HTML from an untrusted source into your application's web UI since doing so would create an exploitable XSS vulnerability. Use an HTML sanitizer instead, which will parse the HTML and conduct runtime checks to make sure each value complies with its contract. Elements that don't adhere to their contract or for which it is difficult to examine the contract in real time are removed by the sanitizer. Because many HTML fragments are unaffected by sanitization, you may use a sanitizer to communicate with other systems that don't use safe types.

Various HTML building libraries focus on various trade-offs between developer efficiency and code readability. However, given that they all uphold the same contract, they ought to all be equally reliable (barring any flaws in their trusted implementations). In fact, at Google, we code-generate the builder functions in a variety of languages from a declarative configuration file to lessen the maintenance effort. The needed contracts for each attribute's value are listed together with the HTML components for each contract. The same configuration file is used by several of

our template systems and HTML cleaners [7], [8]. Closure Templates offers a well-developed open source implementation of safe types for HTML, and work is being done to make type-based security a web standard[9], [10].

## CONCLUSION

This chapter provided a number of design and implementation guidelines for more secure and dependable programming. For sensitive sections of code prone to reliability and security difficulties, such as authentication, authorization, logging, rate limiting, and communication in distributed systems, we specifically advise adopting frameworks as a strong technique. Additionally, frameworks have the tendency to increase developer productivity for both those who create the framework and those who use it and making it much simpler to reason about the code. Aiming for simplicity, using the appropriate tools, employing strong rather than basic types, and regularly sanitizing the codebase are other methods for creating safe and dependable code. It pays off in the long term to make an additional effort to increase security and dependability while building software since doing so will make it easier to examine and repair problems with your programme after it has been launched.

## REFERENCES

- [1] Q. H. Zhu, H. Tang, J. J. Huang, and Y. Hou, "Task Scheduling for Multi-Cloud Computing Subject to Security and Reliability Constraints," *IEEE/CAA J. Autom. Sin.*, 2021, doi: 10.1109/JAS.2021.1003934.
- [2] W. E. Wong and Z. Yang, "Safety, Security, and Reliability of Autonomous Vehicle Software," *Computer (Long. Beach. Calif.)*, 2021, doi: 10.1109/MC.2021.3084655.
- [3] A. Ali *et al.*, "Security, privacy, and reliability in digital healthcare systems using blockchain," *Electron.*, 2021, doi: 10.3390/electronics10162034.
- [4] F. Ahamed, S. Shahrestani, and A. Ginige, "Cloud Computing: Security and Reliability Issues," *Commun. IBIMA*, 2013, doi: 10.5171/2013.655710.
- [5] J. Zhu, Y. Zou, and B. Zheng, "Physical-Layer Security and Reliability Challenges for Industrial Wireless Sensor Networks," *IEEE Access*, 2017, doi: 10.1109/ACCESS.2017.2691003.
- [6] R. Luna and S. A. Islam, "Security and Reliability of Safety-Critical RTOS," *SN Computer Science*. 2021. doi: 10.1007/s42979-021-00753-y.
- [7] Y. Zou, X. Wang, W. Shen, and L. Hanzo, "Security versus reliability analysis of opportunistic relaying," *IEEE Trans. Veh. Technol.*, 2014, doi: 10.1109/TVT.2013.2292903.
- [8] B. Li, X. Qi, K. Huang, Z. Fei, F. Zhou, and R. Q. Hu, "Security-reliability tradeoff analysis for cooperative NOMA in cognitive radio networks," *IEEE Trans. Commun.*, 2019, doi: 10.1109/TCOMM.2018.2873690.
- [9] C. Zhang, J. Ge, F. Gong, F. Jia, and N. Guo, "Security-Reliability Tradeoff for Untrusted

and Selfish Relay-Assisted D2D Communications in Heterogeneous Cellular Networks for IoT,” *IEEE Syst. J.*, 2020, doi: 10.1109/JSYST.2019.2925033.

- [10] D. Koo, Y. Shin, J. Yun, and J. Hur, “Improving security and reliability in Merkle tree-based online data authentication with leakage resilience,” *Appl. Sci.*, 2018, doi: 10.3390/app8122532.

## A STUDY ON TYPES OF TESTING CODES

**Mr. Ahamed Sharif\*\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:ahamedsharif@presidencyuniversity.in

---

### ABSTRACT:

*Site Reliability Engineering (SRE), which assures the dependability, performance, and functioning of software systems, places a strong emphasis on testing code. Rigid testing procedures in SRE enable teams to create high-quality software and preserve system stability by assisting in the identification and prevention of problems before they have an effect on users. This paper gives a general overview of the value of testing codes in SRE and emphasizes important factors for good testing procedures. Unit testing, integration testing, performance testing, and regression testing are just a few of the testing techniques used while evaluating software in SRE. To assure the code's robustness and dependability, these tests evaluate several aspects of it, including its individual components, relationships with other system parts, scalability, and resilience.*

**KEYWORDS:** *Dynamic Program, Fuzz Engine, Integration, Testing.*

---

### INTRODUCTION

An effective system satisfies its defined service level goals, which may also include security assurances, and is robust to faults. Robust software testing and analysis are helpful tools for reducing the likelihood of failure, and they should get special attention throughout the project implementation phase.

Effective testing procedures in SRE give the following advantages:

**Identifying and preventing bugs:** SRE teams may find faults and problems early in the development cycle by doing extensive testing. This lessens the possibility that these problems may appear in production, reducing downtime and user impact.

**Performance Optimization:** Code testing identifies scalability and performance limitations, allowing SRE teams to optimize the code and infrastructure for effective resource use and enhanced system performance.

**System Resilience:** Thorough testing ensures system resilience and lowers the chance of service interruptions by validating the code's capacity to handle unforeseen circumstances, such as excessive traffic or malfunctions.

**Continuous Deployment and Integration:** Continuous integration and deployment are made possible by using automated testing practices, allowing for rapid code updates without affecting system reliability. Automated tests enable quicker release cycles and boost confidence in the quality of code modifications.

**Prevention of Regression:** The prevention of regressions and unwanted side effects in current functionality is made possible by regular testing and regression testing procedures. By doing this, unexpected behavior is avoided and system stability is preserved.

Teams should take the following into account in order to do efficient code testing in SRE:

### **Test Coverage**

To guarantee that crucial code pathways, edge cases, and probable failure situations are fully evaluated, strive for complete test coverage.

### **Test Automation**

Automated testing frameworks and tools should be used to speed up testing, eliminate human error, and enable continuous testing as part of the development process. Include performance testing to assess the system's throughput, response times, and resource utilization under various loads. This helps in locating performance stumbling blocks and improving system performance. For software systems to be reliable, performant, and functional, testing codes in SRE is essential. Effective testing techniques aid in the discovery of problems, performance optimization, system resilience, and regression prevention. SRE teams may produce high-quality software, maintain system stability, and improve user experience by prioritizing rigorous testing and using automated testing frameworks.

We cover a number of testing strategies in this chapter, including unit and integration testing. We also go into depth on other security-related subjects including fuzz testing, static and dynamic programme analysis, and how to make your software more resistant to inputs. Despite their best efforts, the engineers creating your software may inevitably make some errors and fail to consider some edge circumstances. Buffer overflows and cross-site scripting vulnerabilities are two security issues that may be brought on by coding mistakes. Simply said, software may fail in the real world in a variety of ways[1]–[3].

The methods covered in this chapter have distinct cost-benefit profiles and may be used in various phases and environments throughout the software development process. For instance, fuzzing, which involves sending random requests to a system, may assist you with hardening the system's security and dependability. By subjecting the service to a wide range of edge situations, this method may be able to assist you in reducing serving faults and catching information breaches. You'll probably need to carry out extensive upfront testing to find possible vulnerabilities in systems that you can't rapidly and readily repair.

## **DISCUSSION**

### **Unit Testing**

By identifying a variety of flaws in distinct software components before a deployment, unit testing helps improve system security and dependability. This method entails decomposing large software components into smaller, independent "units" that are each tested separately. Code that exercises a specific unit using various inputs chosen by the developer authoring the test makes up unit tests. There are well-known unit test frameworks for many languages, and xUnit-based solutions are widely used. The xUnit paradigm-compliant frameworks enable shared setup and takedown code to run alongside each individual test method. Additionally, these frameworks outline the roles and duties of each component of the testing framework, aiding in the

standardisation of the test result format. Other systems will then have comprehensive information on precisely what went wrong. Popular examples include the built-in unit test package in Python, GoogleTest for C++, go2xunit for Go, and JUnit for Java.

Unit test for a function that checks whether the provided argument is a prime number, written using the GoogleTest framework

```
TEST(IsPrimeTest, Trivial) {  
    EXPECT_FALSE(IsPrime(0));  
    EXPECT_FALSE(IsPrime(1));  
    EXPECT_TRUE(IsPrime(2));  
    EXPECT_TRUE(IsPrime(3));  
}
```

In engineering processes, unit tests are generally performed locally to provide developers quick feedback prior to submitting changes to the source. Unit tests are often executed in CI/CD pipelines prior to a change being merged into a repository's mainline branch. This procedure aims to stop code modifications from altering behavior that other teams depend on.

### **Writing Effective Unit Tests**

The thoroughness and quality of your unit tests will have a big influence on how resilient your programme is. To provide engineers with instant feedback on whether a change has violated intended behavior, unit tests should be quick and reliable. You may make sure that engineers do not break current behavior covered by the relevant tests when they add new features and code by designing and updating unit tests. If a test cannot consistently provide the same findings in an isolated environment, you shouldn't necessarily depend on the test results.

Think of a system that controls the storage bytes that a team is permitted to utilise in a certain datacenter or area. Let's say the system enables teams to ask for more quota when the datacenter has free unallocated bytes. A straightforward unit test may entail verifying quota requests in a collection of fictitious clusters that are only partly populated by fictitious teams and refusing requests that would use more storage than is permitted. Unit tests with a security emphasis may examine how requests with negative bytes are handled or how the code manages capacity overflows for big transfers that produce quota values that are close to the upper bound of the variable types used to represent them. Another unit test may examine if the system responds with the correct error message in the presence of malicious or improper input.

Testing the same code with varied inputs or external factors, like the first beginning quota usages in our case, is often helpful. Unit test frameworks or languages sometimes provide a means to run the same test with varied parameters in order to reduce the amount of duplicated code. By reducing redundant boilerplate code, this strategy may make refactoring tasks less laborious.

### **When to Write Unit Tests**

Developing tests soon after developing the code is a popular tactic for ensuring that the code works as intended. These tests often cover the scenarios that the engineer building the code personally examined. They are generally included in the same commit as the new code. As an

illustration, the storage management application in our example can stipulate that "Only billing administrators for the group that owns the service can request more quota." This kind of requirement may be turned into a number of unit tests [4]–[6].

In companies that use code reviews, a peer reviewer may verify the tests to make sure they're reliable enough to preserve the codebase's quality. A reviewer could see, for instance, that even when new tests are added along with a change, the tests still might pass even if the new code is turned off or inactive. If a reviewer can change a statement in the new code such as `if (condition_1 || condition_2)` to `if (false)` or `if (true)`, and none of the new tests fail, then the test may have missed crucial test cases. See Petrovi and Ivankovi (2018) for additional details on Google's experience automating this kind of mutation testing.

Test-driven development (TDD) approaches encourage engineers to construct unit tests in advance of developing code based on predetermined criteria and anticipated behaviors. Up until the behavior is fully implemented, tests for new features or bug fixes will fail. Engineers go on to the next feature and repeat the procedure once the previous feature has been developed and the tests have passed. In reaction to problem reports or proactive measures to boost system trust, it is typical to gradually integrate and enhance test coverage for existing projects that weren't created using TDD models. But even when you complete full coverage, your project could still have bugs. Even with weakly developed error handling or unknown edge situations, undesirable behavior is still possible.

In response to internal manual testing or code review initiatives, unit tests may also be created. These tests might be created as part of regular development and review procedures or at milestones like a security assessment before a launch. A suggested issue fix's effectiveness and the likelihood that a subsequent refactoring won't generate the same bug may both be confirmed by new unit tests. This kind of testing is crucial when creating access control checks in a system with a complex permission model, for example, or if the code is difficult to reason about and probable defects have an influence on security.

### **How Unit Testing Affects Code**

You may need to create new code with testing facilities or change existing code to make it more testable if you want to increase the thoroughness of your tests. Providing a means to intercept calls to other systems is often part of the refactoring process. You may test code in a number of ways using that introspection capability, for as by making that the code calls the interceptor the right amount of times or with the right inputs. Think about how you would test a piece of code that, when specific criteria are satisfied, opens tickets in a remote issue tracker. Every time the unit test runs, an actual ticket would need to be created, which would make more noise. Even worse, if the issue tracker system is down, this testing technique may randomly fail, defeating the objective of providing timely, accurate test results.

You might take out the direct calls to the issue tracker service and swap them out with an abstraction, like an interface for an `IssueTrackerService` object, to rework this code. When it gets calls like "Create an issue," the testing implementation may collect data. The test may then examine this information to determine if the test passed or failed. The production implementation, in contrast, would call the disclosed API functions when connecting to distant services.

The "flakiness" of a test that relies on real-world systems is drastically reduced by this change. Flaky tests are often more of a hassle than a benefit since they depend on behavior that isn't guaranteed, such as an external dependence or the sequence of components when getting objects from particular container types. In order to prevent developers from developing the habit of disregarding test results while checking in changes, try to resolve problematic tests as they appear.

### **Integration Testing**

Integration testing goes beyond individual units and abstractions, substituting genuine implementations for fictitious or mocked-up versions of abstractions like databases or network services. Integration tests therefore exercise more thorough code pathways. Integration testing may be slower and more unstable than unit testing since you must initialize and establish these other dependencies; also, this method adds real-world factors such network latency when services interact end-to-end during the test execution. A better level of confidence that the system is operating as predicted is the end consequence of switching from testing individual low-level pieces of code to evaluating how they interact when put together.

The intricacy of the dependencies that integration testing addresses determines the many forms that integration testing might take. An integration test may resemble a base class that sets up a few common dependencies (for instance, a database in a predefined state) from which additional tests extend when the dependencies that integration testing requires are very straightforward. Integration tests may become much more complicated as services get more sophisticated, necessitating the coordination of startup or setup of dependencies for the test by supervisory systems. In order to provide standardised integration test setup for common infrastructure services, Google employs teams dedicated just to infrastructure. Depending on the complexity of the codebase and the number of available tests in a project, integration tests may run concurrently with or independently from unit tests in organisations employing a continuous build and delivery system like Jenkins.

### **Writing Effective Integration Tests**

Integration tests, like unit tests, may be impacted by coding design decisions. A unit test mock may only verify that the function was called to submit a ticket to the remote service, continuing our previous example of an issue tracker that creates tickets. A actual client library would be used more often in an integration test. The integration test would interact with a QA endpoint rather than introduce bogus faults in live code. With inputs that prompt calls to the QA instance, test cases would put the application logic to the test. The QA instance might then be queried by supervising logic to confirm that externally visible activities occurred successfully from beginning to finish [7]–[9].

It might take a lot of time and effort to figure out why integration tests fail while all unit tests succeed. You can troubleshoot and figure out where failures occur by using effective logging at crucial logical intersections in your integration tests. Remember that integration tests can only tell you a limited amount about how well those units will conform to your expectations in other situations since they examine interactions between components rather than just individual units, which goes beyond individual units. One of the many reasons why include each kind of testing in your development lifecycle provides value is that one sort of testing often cannot be used in place of another.



### **Dynamic Program Analysis**

Users may do a variety of helpful tasks using programme analysis, such as performance profiling, security-related correctness testing, reporting on code coverage, and eliminating dead code. As will be covered later in this chapter, programme analysis may be done statically to examine software without running it. Here, we emphasize dynamic methods. By executing programs—possibly in virtualized or simulated environments—for reasons other than testing, dynamic programme analysis examines software.

The two most well-known forms of dynamic analysis are performance profilers, which are used to identify performance problems in programs, and code coverage report generators. The dynamic programme analysis tool Valgrind, which offers a virtual machine and numerous tools to parse a binary and determine whether an execution demonstrates many typical defects, was presented in the previous chapter. This section focuses on methods for dynamic analysis that leverage compiler assistance, sometimes known as instrumentation, to find memory-related issues. You may set up instrumentation in compilers and dynamic programme analysis tools to gather runtime statistics on the binaries that the compilers generate, such as performance profiling data, code coverage data, and profile-based optimizations. When the binary is executed, the compiler adds extra instructions and callbacks to a backend runtime library, which exposes and gathers the relevant data. Here, we concentrate on memory usage problems in C/C++ programs that affect security.

### **Fuzz Testing**

A method that supports the testing methodologies stated above is fuzzing, sometimes known as fuzzing testing. Fuzzing is the act of creating a huge number of potential inputs using a fuzz engine (also known as a fuzzer), which are then sent to the target (the code that processes the inputs) through a fuzz driver. The system's response to the input is subsequently examined by the fuzzer. Popular targets for fuzzing include the implementations of file parsers, compression algorithms, network protocols, and audio codecs, which all deal with complex inputs.

Additionally, fuzzing may be used to compare several ways of implementing the same feature. For instance, a fuzzer may produce inputs, send them to both libraries for processing, and then compare the outcomes if you're thinking about switching from library A to library B. Any nonmatching result may be reported by the fuzzer as a "crash," which can assist developers in identifying potential minor behavior changes. As shown in OpenSSL's BigNumfuzzer, this crash-on-different-outputs action is often implemented as a component of the fuzz driver [10].

It is not practical to block every commit on the outcomes of an extensive test since fuzzing may run forever. This implies that a problem may already be checked in when it is discovered by the fuzzer. By producing test cases that engineers would not have thought of, fuzzing serves as a supplement to other testing or analytical techniques that, ideally, would have avoided the problem in the first place. Another unit test may leverage the produced input samples to find faults in the fuzz target, which has the extra advantage of preventing the patch from being regressed by subsequent modifications.

### **How Fuzz Engines Work**

The degree and intricacy of fuzz engines may vary. In order to detect problems, a method known as dumb fuzzing at the low end of the spectrum just takes bytes from a random number

generator and delivers them to the fuzz target. Through connection with compiler toolchains, fuzz engines have become more intelligent. By using the already stated compiler instrumentation capabilities, they may now produce more intriguing and insightful samples. Using as many fuzz engines as you can include into your development toolchain and keeping an eye on metrics like the percentage of code covered are seen to be excellent industry practices. In most cases, it's worthwhile to look into why the fuzzer can't reach further sections if code coverage reaches a plateau.

The specifications or grammars of well-specified protocols, languages, and formats (such HTTP, SQL, and JSON) may provide dictionaries of relevant terms that certain fuzz engines accept. The created parser code may simply reject the input if it includes prohibited keywords, allowing the fuzz engine to generate input that is more likely to be accepted by the programme being tested. Giving a dictionary enhances the chance that fuzzing will find the code you really want to test. Otherwise, you can wind up testing code that only ever detects boring issues and rejects input based on incorrect tokens.

## CONCLUSION

We've just touched the surface of the massive issue of testing software for security and dependability. The testing techniques discussed in this chapter have been crucial in assisting Google teams expand dependably, reducing outages and security issues. They have also helped teams write safe code and remove whole bug classes. It's crucial to include testability into software development from the beginning and to conduct thorough testing at every level of the process. The importance of completely integrating all of these testing and analytic techniques into your engineering processes and CI/CD pipelines is something we want to emphasize at this point. You may find issues more rapidly by combining and using these methods uniformly throughout your codebase. When you launch your apps, which is a subject discussed in the following chapter, your confidence in your ability to find or avoid defects will also increase.

## REFERENCES

- [1] F. C. Ningrum, D. Suherman, S. Aryanti, H. A. Prasetya, and A. Saifudin, "Pengujian Black Box pada Aplikasi Sistem Seleksi Sales Terbaik Menggunakan Teknik Equivalence Partitions," *J. Inform. Univ. Pamulang*, 2019, doi: 10.32493/informatika.v4i4.3782.
- [2] D. R. Farine, "A guide to null models for animal social network analysis," *Methods Ecol. Evol.*, 2017, doi: 10.1111/2041-210X.12772.
- [3] N. Juristo and S. Vegas, "Functional testing, structural testing, and code reading: What fault type do they each detect?," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, 2003, doi: 10.1007/978-3-540-45143-3\_12.
- [4] Y. R. Nasution and M. Furqan, "Aplikasi Mobile Media Pembelajaran Dasar Algoritma dan Pemrograman Berbasis Android," *Syntax J. Softw. Eng. Comput. Sci. Inf. Technol.*, 2020, doi: 10.46576/syntax.v1i1.791.
- [5] S. Molins *et al.*, "Simulation of mineral dissolution at the pore scale with evolving fluid-solid interfaces: review of approaches and benchmark problem set," *Computational Geosciences*. 2021. doi: 10.1007/s10596-019-09903-x.
- [6] A. Zuhair, F. Khadafi, A. M. Andriansyah, B. Saputra, and A. Saifudin, "Teknik

- Pengujian Equivalence Partions untuk Pengujian Aplikasi Sistem Penunjang Keputusan Pegawai Terbaik Menggunakan Black Box,” *J. Teknol. Sist. Inf. dan Apl.*, 2020, doi: 10.32493/jtsi.v3i3.5365.
- [7] S. A. Abdulkareem and A. J. Abboud, “Evaluating Python, C++, JavaScript and Java Programming Languages Based on Software Complexity Calculator (Halstead Metrics),” *IOP Conf. Ser. Mater. Sci. Eng.*, 2021, doi: 10.1088/1757-899x/1076/1/012046.
- [8] M. Nuraminudin, “Analisis Dan Implementasi Onesignal Dalam Pembuatan Aplikasi Mobile Hybrid Lelang Ikan Hias,” *Jurnal.Stmikroyal.Ac.Id*, 2020.
- [9] N. Afzal *et al.*, “Mining peripheral arterial disease cases from narrative clinical notes using natural language processing,” *J. Vasc. Surg.*, 2017, doi: 10.1016/j.jvs.2016.11.031.
- [10] M. L. Hakim, S. Yatmono, A. C. Nugraha, and M. Khairudin, “Autonomous Quadcopter With Image Object Detection Method As A Sender Of Assistance For Covid-19 Patients,” *Int. J. Mechatronics Appl. Mech.*, 2021, Doi: 10.17683/Ijomam/Issue10/V1.2.

## CONCEPT OF DEPLOYING CODES

**Ms. Aashi Agarwal\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA  
Email Id:aashiagarwal@presidencyuniversity.in

---

### ABSTRACT:

*A critical step in Site Reliability Engineering (SRE) is code deployment, which includes the fast and secure introduction of software into real-world settings. It includes a variety of tasks, such as deployment plans, configuration management, version control, and packaging. The important factors and recommended techniques for deploying code in SRE are summarized in this PAPER. For new features, bug patches, or system upgrades to be successfully deployed in SRE, meticulous preparation and collaboration are essential. The abstract emphasizes the value of a clear deployment procedure that minimizes downtime, lowers the chance of mistakes, and ensures system stability.*

**KEYWORDS:** *Configuration, Deployment, Docker Image, Source Code.*

---

### INTRODUCTION

In earlier chapters, we discussed how to write and test your code while keeping security and dependability in mind. That code won't really affect anything, however, until it is created and deployed. As a result, it's crucial to give serious thought to security and dependability at every step of the development and deployment process. It might be difficult to tell whether a deployed artefact is secure just by looking at it. You may feel more secure about the security of a software artefact if there are controls in place at different levels of the software supply chain. Code reviews, for instance, may lower the likelihood of errors and discourage attackers from making nefarious modifications, and automated tests can boost your assurance that the code is working properly. Controls placed around the source, development, and test infrastructure are ineffective if an adversary can deploy directly to your system and get past them. As a result, installations that don't come from the right software supply chain should be rejected by systems. Each stage in the supply chain must be able to demonstrate that it operated correctly in order to achieve this criteria [1]–[3].

Key things to keep in mind while releasing code in SRE are:

#### Deployment Strategies

It is essential to choose the right deployment strategy. Based on the particular needs of the application and the required degree of risk tolerance, options like blue-green deployments, canary releases, or rolling updates should be assessed.

### **Orchestration and automation**

The deployment process may be made more consistent and error-free by automating certain tasks. Infrastructure-as-code techniques and deployment orchestration tools are used to speed up the procedure and guarantee repeatability.

### **Version Control**

Teams are able to monitor changes, handle disagreements, and roll back to earlier versions as needed with the use of effective version control practices, such as utilising source code repositories or version control systems.

### **Testing and Validation**

Before deploying the code, thorough testing and validation methods, including as unit tests, integration tests, and end-to-end tests, should be carried out to ensure its stability and usability. This helps in finding any problems early on in the process.

### **Management of Configurations**

It is possible to customize and maintain software more easily by managing settings independently from the code. Consistent setups across many contexts are ensured through the use of configuration management tools and procedures.

### **Recovery and Rollback**

Planning for potential problems is crucial. Setting up rollback processes and monitoring systems enables speedy recovery in the event of deployment errors or performance problems.

### **Performance Evaluation**

The deployed code is continuously monitored to find performance bottlenecks, scalability problems, or other irregularities. Utilising metrics and monitoring tools makes sure the system achieves performance goals. Organisations may speed the deployment process, reduce interruptions, and preserve the stability and dependability of their systems by adhering to certain principles and best practices.

Careful preparation, automation, and adherence to best practices are required for code deployment in SRE. Smooth transitions are guaranteed by a well conducted deployment procedure, which also lowers risks and keeps software systems' overall dependability high. Organisations may reliably release code into production systems and provide value to their customers while minimizing downtime and interruptions by including these factors into their deployment methods.


## **DISCUSSION**

### **Terminology and Concepts**

We refer to the process of creating, constructing, testing, and deploying a software system as the "software supply chain." The usual duties of a version control system (VCS), a continuous integration (CI) pipeline, and a continuous delivery (CD) pipeline are included in these phases.

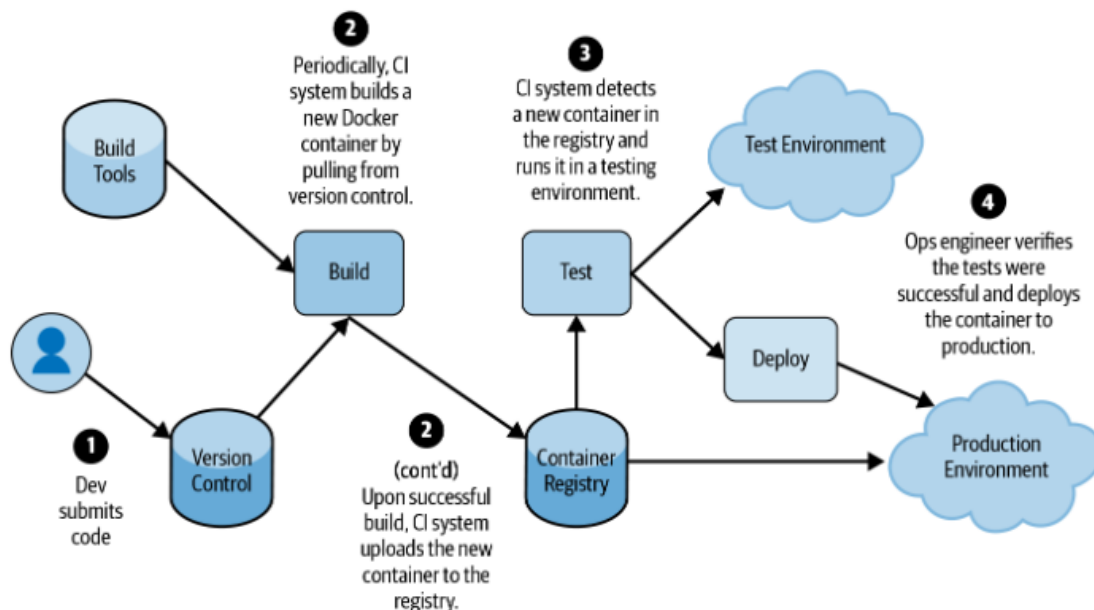
While specifics of implementation differ across businesses and teams, the majority of organisations follow a procedure that resembles Figure 1:

1. A version control system must be checked in with the code.
2. Then, code is created using a checked-in version.
3. The binary must be tested once it is produced.
4. Following deployment, the code is configured and run in a particular environment.



**Figure 1: A high-level view of a typical software supply chain.**

You can typically split your supply chain down into these fundamental components, even if it is more complex than this model. An actual illustration of how a typical deployment pipeline carries out these phases is shown in Figure 2. The software supply chain should be planned to reduce system risks.



**Figure 2: Typical cloud-hosted container-based service deployment.**

Regardless of whether an insider is behaving maliciously, the emphasis of this chapter is on reducing the risks posed by insiders (or malicious attackers posing as insiders), as stated in Chapter 2. An external attacker may try to deploy a backdoored binary using the privileges of a compromised engineer's account, or a well-intentioned engineer could mistakenly build from code that contains unreviewed and unsubmitted modifications. We give equal weight to both circumstances[4]–[6]. The phases of the software supply chain are defined relatively extensively in this chapter. An artefact is any piece of data, such as a file, a package, a Git commit, or a virtual machine (VM) image, and a build is any transformation of input artefacts into output

artefacts. A test is a specific instance of a build in which the output artefact is some logical result typically "pass" or "fail" rather than a file or executable.

It is possible to link builds together and run multiple tests on an artefact. As an example, a release process may "build" binaries from source code, "build" a Docker image from the binaries, and "test" the Docker image by executing it in a development environment. Any assignment of an artefact to an environment is referred to as a deployment. Each of the following may be seen as a deployment:

**Pushing code:**

Sending a command to a server to launch a fresh binary after downloading it

1. Adding a new Docker image to a Kubernetes Deployment object
2. A real or virtual machine's boot process, which loads its basic software or firmware

**Configuration update:**

1. Changing a database structure by use of a SQL statement
2. Changing a command-line flag by updating a Kubernetes Deployment object

Publishing a package or other material that other users will consume:

1. A deb package being uploaded to an apt repository
2. Registering a container registry with a Docker image
3. APK submission to the Google Play Store

**Threat Model**

You must first identify your opponents before securing your software supply chain to reduce hazards. We'll focus on the three categories of enemies listed below during the discussion. Your list of enemies may vary depending on your system and organisation:

1. Benign insiders who are prone to making errors
2. Malicious insiders that attempt to obtain information beyond what their job permits
3. External attackers who get access to one or more insiders' computers or accounts

Next, you need to consider all the ways an opponent may compromise your system by subverting the software supply chain by thinking like an attacker. The frequent hazards listed below are only a few; you should modify this list to reflect the unique dangers facing your organisation. For ease of usage, we refer to engineers as good insiders and to hostile adversaries as both malicious insiders and outside attackers:

- a. A system vulnerability is unintentionally introduced when an engineer proposes a modification.
- b. A malevolent adversary uploads a modification that allows the system to have a backdoor enabled or adds another purposeful vulnerability.
- c. Unreviewed modifications in a locally changed version of the code are mistakenly used when building by an engineer.

- d. A binary with a dangerous configuration is released by an engineer. For instance, debugging capabilities that were originally meant for testing are now enabled in production.
- e. An evil opponent releases a modified binary into production that starts stealing user credentials.
- f. A cloud bucket's ACLs are altered by a hostile attacker, enabling data exfiltration.
- g. The software's integrity key is taken by a malevolent opponent.
- h. An engineer releases an outdated version of the code that has a known bug.
- i. The configuration of the CI system is incorrect since it allows requests to build from any source repository. Consequently, a malevolent opponent may create code from a source repository.
- j. The signature key is exfiltrated by a hostile adversary who uploads a bespoke build script to the CI system. After that, the adversary deploys a malicious binary after signing it with that key[7]–[9].

By tricking the CD system into using a backdoored compiler or build tool, a hostile adversary creates a harmful binary. You may match the risks you discovered to the existing mitigations after you've established a thorough list of prospective adversaries and threats. Any shortcomings in your present mitigation techniques should be documented as well. This activity will provide you a complete view of your system's possible dangers. There is need for improvement with regards to risks that don't have matching mitigations or dangers for which current mitigations have major limits.

### **Ideal Techniques**

You may reduce risks, close any security holes found in your threat model, and constantly enhance the security of your software supply chain by using the recommended practices listed below.

### **Mandating Code Reviews**

Before modifications to the source code are checked in or deployed, it is customary to have a second person (or multiple individuals) examine the changes. Code reviews provide several advantages for a software project in addition to increasing code security: they encourage knowledge exchange and education, establish coding standards, increase code readability, and cut down on errors, all of which contribute to the development of a culture of security and dependability. Code review is a kind of multi-party authorization from a security standpoint, which means that no one has the right to make modifications on their own.

Code reviews must be required for proper implementation. If a competitor can just opt out of the review, they won't be discouraged! Reviews must be thorough enough to identify issues. Any modification must be well understood by the reviewer, along with any consequences it may have for the system. If necessary, the reviewer should approach the author for further information. You may implement required code reviews using a number of publicly accessible technologies. You might set up GitHub, GitLab, or BitBucket to demand a particular amount of approvals for each pull/merge request, for instance. As an alternative, you might combine a source repository set up



to accept just pushes from that review system with standalone review systems like Gerrit or Phabricator.

### **Utilise automation**

The majority of the phases in the software supply chain should ideally be completed by automated systems. There are many benefits to automation. It could provide a reliable, repeatable procedure for developing, evaluating, and deploying software. Removing people from the process helps to decrease labor and error. You fortify the system against subversion by hostile adversaries when you operate the software supply chain automation on a locked-down system.

Imagine an example where engineers create "production" binaries by hand on their workstations as necessary. There are several chances for mistakes to be made in this situation. Engineers may add unreviewed or untested code modifications or unintentionally build from the incorrect version of the source code. In the meanwhile, malevolent adversaries, such as external attackers who have gained access to an engineer's computer, may purposefully replace the locally generated binaries with malicious counterparts. Both of these consequences may be avoided with automation.

### **Verify Not Just People, But Artefacts**

If attackers can get beyond the source, build, and test infrastructure protections by delivering straight to production, their impact is minimal. Verifying who started a deployment is insufficient since that actor might make a mistake or could be deliberately delivering harmful changes. Instead, what is being distributed should be checked in the deployment environments.

Environments for deployment should demand evidence that every automated stage of the deployment process took place. Unless some other mitigating control checks that activity, humans must be unable to circumvent the automation. If you use Google Kubernetes Engine (GKE), for instance, you may utilise Binary Authorization to ensure that only images that have been signed by your CI/CD system are accepted by default and keep an eye on the Kubernetes cluster audit log for alerts when noncompliant images are deployed using the breakglass feature. This method has the drawback of assuming that your whole configuration is safe; for example, it assumes that the CI/CD system only accepts build requests for sources that are permitted in production and that the signing keys (if used) are only accessible by the CI/CD system. A more reliable method of directly evaluating the required qualities with fewer implicit assumptions is described in Advanced Mitigation Strategies.

### **Consider configuration to be code.**

Just as important to security and dependability as the service's code is the configuration of the service. As a result, setup follows all best practices for code versioning and change review. Treat configuration like code by mandating that configuration changes undergo the same review, testing, and checking-in processes as other changes before being deployed. As an example, let's say your frontend server provides a configuration setting that allows you to define the backend. You'd have a serious security and reliability issue if someone pointed your production frontend to a testing version of the backend.

Consider a system that utilizes Kubernetes and keeps its configuration in a version-controlled YAML file as a more real-world example. The deployed configuration is made available by

calling the kubectl programme during the deployment process and passing it the YAML file. It is significantly more difficult to incorrectly setup your service if you limit the deployment process to utilise just "approved" YAML—YAML from version control with necessary peer review.

All the safeguards and recommended practices this chapter suggests for securing the configuration of your service are reusable. These solutions are often considerably simpler to reuse than other ones for safeguarding post-deployment configuration changes, which may need for a whole different multi-party authorization system [10]. Configuration versioning and review are not nearly as common as code versioning and review. Organisations that use configuration-as-code often don't treat configuration with the same rigour as code. Engineers are widely aware that they shouldn't create a production version of a binary from a locally altered copy of the source code, for instance. Those same engineers could release a configuration change without first saving it to version control and requesting approval without giving it a second thought.

Your culture, tooling, and operational procedures must change in order to implement configuration-as-code. You need to prioritize the review process culturally. Technically speaking, you need tools that make it simple to compare proposed changes (like diff and grip) and that let you manually override changes in an emergency.

### **Practical Advice**

Over the years, we've implemented verifiable builds and deployment strategies in many scenarios, and we've learnt a number of lessons. Most of these lessons are more about how to deliver changes that are dependable, simple to debug, and simple to comprehend than they are about the actual technology choices. We hope you will find the suggestions in this part to be helpful.

### **Take It Slowly and Step by Step**

You'll probably need to make a lot of adjustments to be able to provide a highly secure, dependable, and consistent software supply chain, including scripting your build procedures, adopting build provenance, and putting configuration-as-code into practice. The coordination of all of those adjustments could be challenging. The productivity of the engineering department may also be seriously threatened by errors or a lack of functioning in these controls. In the worst instance, a mistake with these settings can result in your service being interrupted. If you concentrate on securing one specific component of the supply chain at a time, you could be more successful. By doing so, you may lessen the chance of a disturbance while also assisting your teammates with acquiring new procedures.

### **Actionable error messages must be provided.**

The error message that appears when a deployment is refused must make it abundantly apparent what went wrong and how to correct it. For instance, if an artefact is denied because its source URI was erroneous, the solution may be to either recreate the artefact from the right URI or to alter the policy to accept that URI. The user should get actionable feedback from your policy decision engine that offers such recommendations. Simply stating that something "does not meet policy" would probably leave the user perplexed and lost.

When creating your architecture and policy language, take into account these user journeys. It might be quite challenging to provide customers useful feedback due to certain design decisions,

therefore look for these issues early. One of our early policy language concepts, for instance, provided a great deal of freedom in the expression of rules but did not allow us to provide actionable error signals. In the end, we gave up on this strategy in favor of a more restricted language that allowed for better error messages.

### Ensure Unambiguous Provenance

In the beginning, the binary provenance was asynchronously uploaded to a database via Google's verified build system. The policy engine then checked the provenance in the database at deployment time using the artifact's hash as a key. While this method mostly functioned as intended, we encountered a serious problem: users may construct the same artefact more than once, leading to several entries for the same hash. Take the empty file as an example. Since many different builds included an empty file in their output, we had literally millions of provenance records linked to the hash of the empty file.

### CONCLUSION

You may strengthen your software supply chain against a variety of insider threats by following the suggestions in this chapter. Automation and code reviews are crucial strategies for reducing errors and raising the cost of an attack for bad actors. These advantages extend to configuration, which has historically received less attention than code. As your organisation expands, you may scale using artifact-based deployment rules, especially those that include binary provenance and verified builds. They also provide defence against knowledgeable attackers. These suggestions work together to make sure that the code you created and tested is the code that is really used in production. Nevertheless, despite your best efforts, it's likely that your code won't always operate as intended. When that occurs, you may make use of some of the debugging techniques covered in the next chapter.

### REFERENCES

- [1] J. A. Barriga, P. J. Clemente, E. Sosa-Sanchez, and A. E. Prieto, "SimulateIoT: Domain Specific Language to Design, Code Generation and Execute IoT Simulation Environments," *IEEE Access*, 2021, doi: 10.1109/ACCESS.2021.3092528.
- [2] B. Sieglin *et al.*, "Exploitation of DevOps concepts for the ASDEX Upgrade DCS," *Fusion Eng. Des.*, 2021, doi: 10.1016/j.fusengdes.2021.112332.
- [3] J. H. Xue, "Algorithmic Vulnerability in Deploying Vaccination Certificates in the European Union and China," *Eur. J. Risk Regul.*, 2021, doi: 10.1017/err.2021.32.
- [4] I. A. Elgendy, W. Z. Zhang, H. He, B. B. Gupta, and A. A. Abd El-Latif, "Joint computation offloading and task caching for multi-user and multi-task MEC systems: reinforcement learning-based algorithms," *Wirel. Networks*, 2021, doi: 10.1007/s11276-021-02554-w.
- [5] M. J. J. Douglass, "Book Review: Hands-on Machine Learning with Scikit-Learn, Keras, and Tensorflow, 2nd edition by Aurélien Géron," *Phys. Eng. Sci. Med.*, 2020, doi: 10.1007/s13246-020-00913-z.
- [6] A. Henley, J. Ball, B. Klein, A. Rutter, and D. Lee, "An Inquisitive Code Editor for Addressing Novice Programmers' Misconceptions of Program Behavior," in *Proceedings*

- *International Conference on Software Engineering*, 2021. doi: 10.1109/ICSE-SEET52601.2021.00026.
- [7] Aurélien Géron, *Hands-on machine learning with Scikit-Learn, Keras and TensorFlow: concepts, tools, and techniques to build intelligent systems*. 2019.
- [8] X. Liu, M. Xu, T. Teng, G. Huang, and H. Mei, "MUIT: A Domain-Specific Language and its Middleware for Adaptive Mobile Web-Based User Interfaces in WS-BPEL," *IEEE Trans. Serv. Comput.*, 2019, doi: 10.1109/TSC.2016.2633535.
- [9] R. Inam, J. Carlson, M. Sjödin, and J. Kunčar, "Predictable integration and reuse of executable real-time components," *J. Syst. Softw.*, 2014, doi: 10.1016/j.jss.2013.12.040.
- [10] S. Newman, *Building Microservices*. 2015.

---

## INVESTIGATING SYSTEMS AND SECURE DEBUGGING ACCESS

**Dr. Topraladoddi Madhavi\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA,  
Email Id:madhavit@presidencyuniversity.in

---

### ABSTRACT:

*Site Reliability Engineering (SRE), which entails diagnosing and fixing faults in production settings while assuring the security and integrity of the system, places a high priority on looking into systems and securing debugging access. The main points and recommended methods for analysing systems and setting up safe debugging access in SRE are summarized in this abstract. To find the underlying causes of events, mistakes, or performance problems, system investigations in SRE must be rigorous and data-driven. The abstract emphasizes how crucial efficient monitoring, recording, and incident response protocols are to enabling prompt problem investigation and resolution. SRE teams may successfully analyse system faults, pinpoint root causes, and put remediation plans into place while protecting the security and integrity of the system by paying attention to these factors and putting best practices into practice.*

**KEYWORDS:** Access Control, Compliance, Debugging, Monitoring.

---

### INTRODUCTION

Most systems inevitably collapse. You need the right knowledge in addition to having access to sufficient logs and information sources for debugging in order to effectively analyse a complicated system. Additionally, you need to consider security and access control while designing your logging systems. In this chapter, we go through debugging methods and provide a few tips on what to do if you get stuck. The contrasts between troubleshooting a system problem and looking into a security issue are then covered, along with considerations to consider when determining which logs to keep. Finally, we consider how to maintain the security and dependability of these important informational resources. In SRE, secure debugging access and system investigation are crucial factors to take into account.

**Monitoring and Alerting:** By putting in place thorough monitoring and alerting systems, SRE teams are better able to proactively identify system abnormalities and react to events swiftly. Setting proper criteria and configuring pertinent alerts enables prompt inquiry and resolution.

**Logging and tracing:** Appropriate systems for logging and tracing provide important insights into the activities and interactions of different system components. By allowing SRE teams to examine log data and track requests across several services, the use of distributed tracing and centralized logging technologies promotes effective investigation.

**Post-Mortems and Incident Response:** After an incident, doing a post-mortem study and establishing clearly defined incident response procedures may assist determine the underlying

reasons and build preventative measures. It promotes a culture of learning and continual development to conduct blameless post-mortems.

**Access Controls and Permissions:** Restricting authorized personnel's access to production environments and debugging tools helps to maintain the system's security and integrity. Only those with the required access rights may access sensitive data or carry out debugging tasks thanks to the implementation of role-based access control (RBAC) and least privilege principles.

**Secure Debugging Methods:** Using secure debugging methods, such as secure shell (SSH) tunnels and encrypted communication channels, may help safeguard sensitive data during debugging sessions. An additional layer of protection is added by using isolated debugging environments or secure remote debugging tools.

**Compliance and auditing:** Monitoring access logs, debugging activities, and industry laws aid in maintaining accountability and spotting any unauthorized access or abuse of debugging tools.

As a result, proactive monitoring, efficient incident response procedures, access restrictions, and secure debugging methods must all be used in conjunction with SRE to investigate systems and secure debugging access. Organisations may swiftly identify and fix problems, lessen the effect of accidents, and protect the security and dependability of their systems by integrating these factors into their practices.

In a perfect world, everyone would create flawless systems, and everyone would utilise them with the best of intentions. In practice, you'll run across bugs and have to look into security. You may see opportunities for improvement and locations where you can simplify and optimize processes as you watch a system operate in production over time. All of these jobs call for suitable system access, debugging, and investigation methods. However, there is a chance that this access might be exploited if even read-only debugging access is granted. You need to put in place suitable security measures to mitigate this danger. Additionally, you must carefully balance the security requirements for storing and accessing sensitive data with the debugging needs of developers and operational workers[1]–[3].

## DISCUSSION

### From Debugging to Investigation

Debugging is not popular. Bugs appear when least expected. When a defect will be resolved or a system will be "good enough" for widespread usage, it may be difficult to predict when. Most individuals find writing fresh code more enjoyable than debugging older programs. Debugging could seem like an unfulfilling task. However, it's essential, and if you look at the practice through the lens of gaining new knowledge and skills, you could even find it pleasurable. Debugging, in our opinion, also improves our programming skills and serves as a reminder that we are not always as brilliant as we think we are.

### For instance, temporary files

Take a look at the issue we (the writers) fixed two years ago. When we got a warning that a Spanner database was about to run out of storage, we started looking into it. As we proceeded through the debugging process, we posed the following inquiries to ourselves:

#### a. What led to the database's storage space being exhausted?

A quick triage revealed that the issue was brought on by a buildup of many little files being produced in Google's enormous distributed file system, Colossus, which was probably brought on by a shift in user request volume.

**b. What was responsible for all the little files?**

The files were a consequence of the Spanner server having insufficient memory, according to service metrics that we looked at. Recent writes (updates) were normally cached in memory; however, when the server ran out of memory, it flushed the data to files on Colossus. Unfortunately, there wasn't enough RAM on each server in the Spanner zone to provide updates. As a consequence, each server flushed a high number of little files to Colossus as opposed to a reasonable quantity of bigger, compressed files.

**c. What was being utilised from the memory?**

Every server used a Borg job to run (in a container), which limited the amount of RAM it could use. We simply ran the slabtop command on the production computer to see where in the kernel RAM was consumed. We find that the cache for directory entries (dentry) was the main memory consumer.

**d. The dentry cache was really large, why?**

For each flush operation, we estimated that the Spanner database server was producing and removing a small number of temporary files. The issue became worse as the dentry cache grew bigger with each flush operation.

**e. How might we prove our theory?**

We developed and ran a programme on Borg to create and delete files repeatedly in order to replicate the problem in order to verify this notion. The dentry cache had used all of the RAM in its container after a few million files, supporting the theory.

**f. Was there a kernel flaw here?**

We investigated the anticipated behavior of the Linux kernel and discovered that it caches the absence of files, a characteristic that certain build systems need to guarantee acceptable speed. When the container is full, the kernel normally removes items from the dentry cache. However, since the Spanner server flushed updates frequently, the container was never sufficiently full to start evictions. We solved this problem by stating that caching of temporary files was not required.

Many of the ideas we cover in this chapter are illustrated by the debugging process as it is detailed here. The most crucial lesson to learn from this tale is that we fixed the problem, and you can too! It didn't take any magic to solve and resolve the issue; it only needed careful, methodical research. To clarify the features of our investigation:

- a. We investigated the issue utilising the system's existing logs and monitoring infrastructure when it started to degrade.
- b. Even though it happened in code that the debuggers had never seen before in kernel space, we were still able to track down the problem.

- c. Despite the fact that the problem had probably been for a while, we had never seen it prior to this outage.
- d. The system was not broken in any way. Every component was operating as anticipated.
- e. The fact that temporary files might continue to use memory after being erased startled the Spanner server's creators.
- f. Using tools made available by the kernel developers, we were able to troubleshoot the kernel's memory use. We were trained and experienced in debugging procedures, so even though we had never used these tools previously, we were able to advance very rapidly.
- g. The problem was first misinterpreted as a human mistake. We only modified our thoughts after carefully considering our facts.
- h. Before making improvements to the system, we identified the underlying issue by formulating a hypothesis and then devising a method to verify it.

### **Separate horses from zebras.**

Do you immediately picture horses or zebras when you hear hoofbeats? Teachers sometimes ask this question of medical students who are learning how to triage and identify illnesses. It serves as a reminder that most illnesses are widespread, with horses—not zebras—causing the majority of hoofbeats. You can see why a medical student would find this advice useful: they don't want to believe a set of symptoms indicates a rare illness when, in reality, the ailment is common and easily treatable.

On the other hand, skilled engineers will notice both typical and uncommon occurrences when the scale is great enough. The goal of computer system developers should be to entirely eradicate all issues. Rare problems become more prevalent as a system becomes bigger and its operators learn to solve common issues. Bryan Cantrill once said, "Over time, the horses are found; only the zebras are left."

Think about the very uncommon problem of bit flip memory corruption. There is a less than 1% probability each year that an error-correcting memory module would have an uncorrectable bit flip that will cause a system crash. It's doubtful that an engineer troubleshooting a sudden crash would say, "I bet this was caused by an incredibly unlikely electrical malfunction in the memory chips!" On a very big scale, nevertheless, these rarities become commonplace. A fictitious cloud business with 25,000 computers may utilize 400,000 RAM chips for memory. Given the probability of a 0.1% annual risk of irreparable faults per chip, the service's scope may result in 400 instances per year. The cloud service's users will probably notice a memory failure every day [4]–[6].

Even though debugging these uncommon occurrences might be difficult, it is possible with the correct data. For instance, Google hardware experts once discovered that certain RAM chips failed considerably more often than they should have. They were able to trace the faulty DIMMs (memory modules) down to a single supplier because to asset data that helped them find the source of the problems. The engineers discovered the underlying problem after a protracted debugging and investigation: an environmental breakdown in a clean room, at a single plant where the DIMMs were manufactured. The issue was a "zebra" a uncommon bug that can only be seen at scale.



### **Set aside time for inquiry and debugging**

Debugging often requires a lot of time many continuous hours of work as do security investigations (to be covered later). Debugging for the temporary files case in the preceding section took between 5 and 10 hours. When managing a significant issue, isolate debuggers and investigators from the minute-by-minute reaction to allow them room to concentrate. Debugging promotes patient, systematic, persistent methods where employees are ready to go the extra mile and double-check their work and premises. The issue with temporary files also provides a bad illustration of debugging since the original responder initially identified the outage as being brought on by user traffic and attributed users' poor system behaviour to it. The crew was now dealing with operational overload and pager weariness brought on by non-urgent pages.

Following your observations, note your expectations and the reasons for them. Between your conceptual picture of the system and its real implementation, bugs often hide. In the case of the temporary files, the programmers believed that removing a file also deleted all references to it.

### **Recognise what is typical for your system.**

Debugging software often begins with anticipated system behavior. Here are a few instances based on our knowledge:

1. Near the conclusion of its shutdown code, a binary named abort. The fascinating failure was really the cause of the call to shut down, but new devs didn't realize this until they began debugging the call after seeing the abort call in the logs.
2. To check if the network is improperly meddling with DNS, the Chrome web browser launches three random domains (like cegzaukxwefark.local) and tries to resolve them. Even the Google investigative team mistook these DNS resolves for malicious software attempting to resolve a hostname for a command-and-control server.

Even if the events seem pertinent or suspicious, debuggers often need to filter them away. The constant presence of background noise and active opponents who could be attempting to conceal their activities provide an additional challenge for security investigators. Before you can see more significant problems, you often need to filter out normal noisy activities like automatic SSH login brute forcing, authentication difficulties brought on by users' incorrect password input, and port scanning.

Establishing a baseline of system behavior when no issues are suspected is one technique to comprehend typical system behavior. If you already have a problem, looking at previous logs from before the issue started could help you determine your baseline.

For illustration, we discussed a widespread YouTube outage brought on by a modification to a general logging library. The modification resulted in the servers' failure due to memory exhaustion (OOM). Our post-outage analysis questioned if the outage had an impact on the quantity of OOMs for all other Borg processes due to the library's widespread use inside Google. Although the logs indicated that we had several OOM circumstances on that particular day, a comparison of the data with data from the preceding two weeks revealed that Google experiences numerous OOM conditions on a daily basis. Despite being a major problem, it had no appreciable impact on the OOM statistic for Borg jobs.

Avoid normalising deviation from recommended practices. Over time, defects often take on "normal behavior" and you stop seeing them. For instance, we previously worked on a server that had heap fragmentation use around 10% of its RAM. After claiming for many years that a loss of 10% was predicted and hence acceptable, we looked at a fragmentation profile and identified significant potential to save RAM. You can develop a blind spot as a result of operational stress and alertness fatigue, which would normalize deviation. The process of writing documentation and describing a system to others can also make you wonder how well you understand it. To address normalised deviance, we actively listen to newcomers to the team, and to promote fresh perspectives, we rotate people in and out of on-call rotations and response teams. In addition, we assess our blind spots using Red Teams.

### **Secure and robust debugging access**

You often require access to the systems and the data they hold in order to fix problems. Can an infected or malicious debugger see sensitive data? Can a security system failure be fixed (remember: all systems fail!)? You must guarantee the dependability and security of your debugging systems [7], [8].

### **Reliability**

Another method that systems might fail is via logging. For instance, a system could run out of disc space where logs are stored. In this case, failing open involves still another trade-off: although the method may increase system resilience overall, it also increases the risk of your logging mechanism being interfered with by an attacker. Consider circumstances in which you may need to troubleshoot or fix the security systems themselves. Make the required trade-offs to avoid locking oneself out of a system while maintaining its security. In this situation, you may think about retaining a set of emergency-only credentials that, when used, trigger high-confidence alerts offline and in a safe area. For instance, a recent network disruption at Google resulted in significant packet loss. The authentication system failed when responders tried to retrieve internal credentials because it was unable to contact one backend. However, the rescuers were able to access and repair the network thanks to emergency credentials.

Administrators may assume the identity of a user and see the user interface (UI) from their point of view on the Security One phone support system we worked on. This approach was excellent for debugging since it allowed you to easily and rapidly duplicate a user's issue. However, there is room for misuse in this kind of system. Securing debugging endpoints is necessary, from impersonation to bare-bones database access.

Debugging odd system behaviour often doesn't need access to user data. For instance, the rate and calibre of the bytes travelling across the cable may often be used to identify difficulties with TCP communication. Data in transit may be shielded from any potential surveillance attempts by other parties by using encryption. As a fortuitous byproduct, additional engineers will now have ready access to packet dumps as required. One error that may be made is to disregard the sensitivity of metadata. By observing associated access habits, such as a person visiting a dating site and a divorce lawyer in the same session, a hostile actor may still learn a lot about a user via metadata. The dangers associated with treating metadata as nonsensitive should be carefully considered. Additionally, certain analyses do call for the use of real data; for instance, identifying frequently accessed entries in a database and then determining the reasons behind those accesses. A low-level storage issue resulting from a single account getting tens of thousands of emails per

hour was previously resolved by our team. More details on access control in these circumstances may be found in Zero Trust Networking [9], [10].

## CONCLUSION

Iterate by considering some previous investigations and determining what details may have aided in the investigation or debugging of a problem. As a process of continual improvement, debugging involves adding new data sources and seeking out methods to increase observability. Safety-conscious design. You need logs. System and data storage access is necessary for debugging. However, when your data storage grows, both logs and debugging endpoints may attract the attention of malicious parties. Create logging systems that will capture the data you need while simultaneously imposing strict rights, privileges, and regulations to access that data. Even the finest debuggers are sometimes left in the dark, since both debugging and security investigations often rely on lucky deduction and flashes of insight. Keep in mind that preparation pays off; by having logs and a method for categorizing and studying them available, you can take advantage of opportunities as they present themselves.

## REFERENCES

- [1] R. L. Baskerville, "Investigating Information Systems with Action Research," *Commun. Assoc. Inf. Syst.*, 1999, doi: 10.17705/1cais.00219.
- [2] A. Yaqoob, "A Critical Analysis of Crime Investigating System in India," *Int. J. Trend Sci. Res. Dev.*, 2019, doi: 10.31142/ijtsrd23675.
- [3] M. D. Myers, "Investigating Information Systems with Ethnographic Research," *Commun. Assoc. Inf. Syst.*, 1999, doi: 10.17705/1cais.00223.
- [4] J. Lane, "Intersectionality and Investigating Systems of Privilege and Oppression in Nursing and Health Research," *The Canadian journal of nursing research = Revue canadienne de recherche en sciences infirmieres*. 2020. doi: 10.1177/0844562119900741.
- [5] A. Al-Hunaiyyan, R. Alhajri, B. Alghannam, and A. Al-Shaher, "Student Information System: Investigating User Experience (UX)," *Int. J. Adv. Comput. Sci. Appl.*, 2021, doi: 10.14569/IJACSA.2021.0120210.
- [6] G. Paré, "Investigating Information Systems with Positivist Case Research," *Commun. Assoc. Inf. Syst.*, 2004, doi: 10.17705/1cais.01318.
- [7] R. L. Baskerville, "Investigating Information Systems with Action Tutorial Investigating Information Systems with Action," *October*, 1999.
- [8] N. Savona, C. Knai, and T. Macauley, "Investigating system-level drivers of obesity with adolescents: a group model-building exercise," *Lancet*, 2019, doi: 10.1016/s0140-6736(19)32880-6.
- [9] M. E. Popescu, E. Militaru, A. Cristescu, M. D. Vasilescu, and M. M. M. Matei, "Investigating health systems in the European Union: Outcomes and fiscal sustainability," *Sustain.*, 2018, doi: 10.3390/su10093186.
- [10] A. Zutz *et al.*, "A dual-reporter system for investigating and optimizing protein translation and folding in *E. coli*," *Nat. Commun.*, 2021, doi: 10.1038/s41467-021-26337-1.

---

**BUILDING A CULTURE OF SECURITY AND RELIABILITY****Dr. Shrishail Anadinni\***

\*Assistant Professor,  
Department Of Civil Engineering,  
Presidency University, Bangalore, INDIA,  
Email Id:shrishail@presidencyuniversity.in

---

**ABSTRACT:**

*A crucial component of Site Reliability Engineering (SRE) is creating a culture of security and dependability that encourages an organization's employees to be proactive and resilient. The main ideas and tactics for creating an SRE culture that places a high priority on security and dependability are summarized in this paper. The protection and stability of systems and data must be shared by all levels of the organisation in order to foster a culture of security and dependability. In developing such a culture, the abstract emphasizes the significance of leadership support, education and training, teamwork, and constant progress. A comprehensive strategy that includes leadership support, education and training, cooperation, empowerment, accountability, continuous improvement, and the use of metrics is needed to create a culture of security and dependability in SRE. Organisations may build an atmosphere where security and dependability are engrained in every element of their operations, resulting in robust and reliable systems, by developing and fostering such a culture.*

**KEYWORDS:** *Dependability, Education Training, Reliability, Security.*

---

**INTRODUCTION**

A healthy culture of security and dependability is described in this chapter in detail. We also discuss how, when it comes time to make changes, you may influence organizational culture by making wise decisions. Finally, we provide advice on how to persuade teams throughout the organisation and leadership to support security and dependability. Building a culture of security and dependability in SRE involves many important factors, including:

**Leadership Assistance**

It is essential for the leadership to support security and dependability measures. The building blocks for the culture's success include clear expectations, resources, and setting an excellent example.

**Training and Education**

Employees may get the skills and information needed to contribute to a safe and dependable environment by receiving frequent education and training programs on security best practices, incident response, and reliability engineering. This encourages a culture of ongoing learning and development.

**Communication and Cooperation**

Cross-functional cooperation and open lines of communication between the development, operations, and security teams provide a common concept of security and dependability. This encourages information exchange, pro-active problem-solving, and successful incident response.

### **Responsibility and Empowerment**

Enabling teams to take charge of security and dependability by giving them the freedom and tools they need to put best practices into practice. Teams should be held responsible for achieving security and dependability objectives to underline the significance of these concepts.

### **Constant Development**

Continuous improvement is facilitated by establishing feedback loops, performing post-mortems, and applying lessons discovered from mishaps and vulnerabilities. Supporting a proactive and resilient mentality involves fostering a blameless culture where errors are seen as teaching moments.

### **Metrics for security and dependability**

The success of security and reliability practices may be evaluated by defining and monitoring key indicators for security incidents, system uptime, and reliability. Utilising metrics to track development and convey the results of security and reliability activities helps the organisation understand how important they are.

Organisations may promote a culture that values security and dependability by adhering to these recommendations and putting initiatives into place. A strong culture of security and dependability not only reduces risks and safeguards priceless assets but also fosters trust among customers and stakeholders, eventually enhancing the organization's long-term performance. We've included some procedures below that we've found to be helpful at Google and elsewhere, but bear in mind that no two businesses are alike, so you'll need to modify these procedures to fit the culture of your business. Additionally, you'll probably discover that not all of these tactics may be used. This chapter serves as a manual and a point of reference for ongoing thought. It's important to note that at Google, we strive to make improvements in each strategy as part of creating our entire culture, rather than putting the advice we give here into flawless practice every day.

Organisations that see the value of effective security and dependability create a culture centered on these principles. By making culture a team effort—the responsibility of everyone, from the CEO and their leadership team to technical leaders and managers to the people who design, implement, and maintain systems organizations that explicitly design, implement, and maintain the culture they seek to embody succeed [1]–[3].

Consider the following scenario: Just last week, the CEO informed your whole organisation that closing the following Big Deal was essential to the company's future. You discovered evidence of an attacker on the company's systems this afternoon, and you are aware that these systems must be taken down. Customers will be upset, and the Big Deal might be in jeopardy. Additionally, you are aware that while many employees were on vacation and everyone was working under a tight schedule for the Big Deal, your team may be held accountable for not implementing security fixes last month. What choices does your workplace culture encourage

workers to make in this circumstance? Despite the possibility of postponing the Big Deal, a healthy organisation with a solid security culture would urge staff to report issues right away.

Imagine the frontend development team unintentionally pushes a large update meant for staging to the live production system while you're busy looking into the malicious intruder. Customers are overflowing the support hotline as a result of the glitch, which knocks the company's income stream down for more than an hour. Your clientele's trust is fast deteriorating. Employees would be encouraged to rethink the procedure that permitted an unintentional frontend push in an environment where dependability is valued so that teams could balance the demands of consumers with the potential for missing or postponing the Big Deal.

In these circumstances, cultural norms need to promote blameless postmortems in order to identify patterns of failure that may be corrected, preventing dangerous conditions from occurring in the future.<sup>1</sup> Companies with strong cultures are aware that although being hacked once is bad, it is considerably worse to be hacked again. They are also aware that 100% is never the ideal reliability aim and that, by balancing reliability and velocity, they may keep users happy by using techniques like error budgets<sup>2</sup> and safe code push restrictions. Finally, organisations with a strong security and dependability culture are aware that users value openness when issues unavoidably happen and that concealing such accidents may reduce user confidence.

In order to create a culture of security and dependability, this chapter discusses several patterns and anti-patterns. Although we believe that this material will be useful to businesses of all sizes, culture is a distinctive aspect of an organisation that is developed in light of its specific difficulties and characteristics. The cultures of different organisations may differ, and not all of the tips we provide here may be universally relevant. It's doubtful that any organisation will be able to implement all of the practices we examine, but this chapter is aimed to provide a variety of perspectives on the subject of culture. The slightly idealized perspective we provide here won't be entirely useful in realistic applications. Even at Google, we don't always get culture right, so we're always looking for ways to make it better. We hope you'll discover some of the many perspectives and choices given here that could fit your surroundings.

## DISCUSSION

### **A Healthy Security and Reliability Culture: Definition**

A strong team culture may be intentionally developed, put into place, and kept up, much like healthy systems. We have concentrated on the technical and procedural aspects of creating healthy systems throughout this book. There are other design considerations for creating wholesome cultures. In actuality, building, deploying, and sustaining safe and dependable systems depend heavily on culture.

### **Default security and dependability culture**

It might be easy to put off thinking about security and dependability until later in a project's lifespan, as we cover in Chapter 4. Although this delay seems to improve initial velocity, it actually reduces sustained velocity and may raise the cost of retrofits. These retrofits may fail if they are implemented inconsistently or over time build technical debt. Imagine having to independently find a seat belt seller, someone to assess the safety of the windscreen, and an inspector to verify the airbags when purchasing a vehicle to demonstrate this point. Resolving

dependability and safety issues only after a vehicle has been built would put a heavy responsibility on the customer, who may not be in the greatest position to judge whether the remedies put in place are enough. Additionally, it can result in different manufacturing processes for every two automobiles.

This comparison highlights the need for systems to be trustworthy and secure by design. Consistency is simpler to maintain when security and reliability decisions are established throughout a project's lifespan. Additionally, as they are incorporated into the system, they may lose their visibility to the user. Reverting to the automobile example, customers may believe that they act morally without giving safety features like seat belts, windscreens, or rear-view cameras any attention. Employees are encouraged to talk about these issues early in the project lifecycle—for instance, at the design stage and during each iteration of execution in organisations with strong cultures of security and dependability by default. According to the software development lifecycle, products' security and dependability will naturally improve as they become older.

People who are building, managing, and implementing systems will find it simpler to automatically and transparently include security and dependability considerations. For instance, you may implement automation for testing, sanitizers, continuous builds, and vulnerability finding. Common vulnerabilities like XSS and SQL injection may be avoided by developers with the aid of application frameworks and common libraries. Memory corruption issues may be avoided with advice on selecting the suitable programming languages or programming language features. This kind of automated security should be fairly obvious to developers and attempts to lower friction (such as slow code audits) and mistakes (such as flaws not found during review). Employees should grow to trust these implementations as systems become better at following these security and reliability standards[4]–[6].

### **Awareness Culture**

Members of an organisation may effectively produce positive results when they are aware of their security and dependability duties and are aware of how to fulfil them. For instance, an engineer who accesses critical systems would need to take more precautions to keep their account safe. People who often communicate with outside parties as part of their jobs can encounter more phishing emails. When an executive visits specific regions of the globe, their risk may increase. Healthy cultures encourage knowledge of these issues and promote them via educational initiatives.

Building a good security culture requires using awareness and education tactics. In order to engage students with the subject matter, these efforts should aim to be lighter and enjoyable. Depending on the information's presentation, level of prior knowledge, and even individualized characteristics like age and background, people remember various sorts of information at various rates. In our experience, interactive learning techniques like hands-on laboratories result in a greater percentage of student retention than passive learning techniques like viewing videos. When raising awareness, think carefully what kinds of information you want people to remember and how you want them to learn in order to optimize for the optimal learning experience. Google employs a variety of strategies to educate staff members about security and dependability. We generally need all staff to attend yearly training. Then, we reaffirm these signals using specialized programs for certain purposes. Here are several strategies that, after years of using these programs at Google, we've found to be effective:

### Interactive talks

Talks that include audience involvement may be an effective approach to convey difficult material. For instance, discussing the most important root causes and mitigations for big security and reliability problems has made it clearer to Google workers why we prioritize these issues. These participatory talks, according to our research, also motivate individuals to bring up any problems they come across, whether they bugs in the code that might bring down systems or suspicious activities on their workstations. This practice gives employees a sense of belonging to the group that strengthens the organization's security and dependability.

### Games

Another strategy to raise awareness is to include security and dependability into games. bigger organisations may be better equipped to provide participants with flexibility over when they take the training and the chance to retake it if they so want since these strategies have a tendency to scale more efficiently to bigger organisations. Our XSS game (shown in Figure 1), which we developed, has been quite effective at educating developers about this widespread online application vulnerability.



Figure 1: A security training game.

### Reference Documentation

We've discovered that it's crucial to provide developers solid documentation they can turn to when necessary, even if reading documentation may have a lower retention rate for learning than approaches like hands-on activities. Reference documentation is crucial since it's difficult to remember all the subtleties of security and dependability at once. Google has a list of internal security best practices that engineers may search for solutions to issues as they emerge for help



on typical security issues. All documentation must be clearly owned, maintained current, and deprecated when it is no longer necessary.

### **Awareness Campaigns**

Notifying developers of current security and reliability concerns and advancements might be challenging. Google offers weekly technical advice in a one-page style as a solution. These "Testing on the Toilet" episodes are made available in all Google workplaces' toilets. The program's primary focus was on improving testing, although it also sometimes covers security and reliability-related problems. A flyer placed in a conspicuous area is a useful technique to convey advice and offer inspiration[7]–[9].

### **Problem-solving and Escalations**

Despite best efforts, there are instances when the need to decide on a security or dependability modification might suddenly and explosively appear. Perhaps a significant outage or security incident indicates that you urgently need extra people and resources. Or maybe the decision-making process isn't functioning because two teams have different ideas about how to tackle an issue. You may need to go up the management chain to find a solution in these kinds of circumstances. We advise using the following rules to handle escalations:

Get a group of coworkers, mentors, tech leaders, or managers together to discuss the matter and provide their perspectives from both sides. Before opting to escalate, it's typically a good idea to go through the matter with an objective third party. Ask the group to summarize the circumstance and suggested management decision possibilities. Make this summary as succinct as you can. Mention any pertinent supporting facts, chats, bugs, designs, etc., and keep your tone absolutely factual. Make as much sense as you can about the probable effects of each choice.

To secure additional agreement on potential solutions, distribute the summary to the team leadership on your own team. For instance, it can be necessary to escalate many situations at once. You may wish to combine escalations or highlight different elements of related circumstances. Plan a meeting to inform all concerned management chains of the problem and name the proper decision-makers in each chain. The decision-makers should then meet separately to debate the matter or reach a formal decision [10].

For instance, when the product team and the security reviewer cannot agree on the appropriate course of action, security concerns at Google may need to be escalated. The security staff starts an escalation in this situation. The two senior executives in the two organisations then decide whether to reach a compromise or execute one of the solutions recommended by the security team or the product team. We don't see these escalations as confrontational since we incorporate them into our everyday workplace culture.

### **CONCLUSION**

You may develop, implement, and sustain an organization's culture over time to achieve security and dependability objectives, just as you can do the same with technology. Engineering efforts should be carefully weighed against reliability and security initiatives. There are significant cultural aspects of engineering that, when combined or even considered separately, may lead to more reliable systems. Improvements in security and dependability might cause anxiety or worry about escalating conflict. There are ways to allay these concerns and encourage support from

those affected by the changes. The trick is to make sure your objectives are well-aligned with stakeholders, especially leadership. People may be persuaded to embrace change more quickly by empathizing with them and putting a strong emphasis on usability. You could have more success persuading others that your adjustments are sensible if you spend a little time considering how they see change. Since no two cultures are alike, as we said at the introduction of this chapter, you will need to modify our suggested techniques for your own company. As a result, you will probably not be able to use all of these tactics. It could be helpful to choose the areas that require the greatest development for your organisation and work on them over time, following Google's long-term strategy of ongoing improvement.

## REFERENCES

- [1] O. Ugizaqiah, F. Supriani, and M. Islam, "Analisis Faktor-Faktor Yang Mempengaruhi Masyarakat Dalam Pengembangan Bangunan Rumah Tipe (Studi Kasus Perumahan Surabaya Permai 4)," *Inersia J. Tek. Sipil*, 2020, doi: 10.33369/ijts.12.2.75-82.
- [2] D. J. Kim, "Self-perception-based versus transference-based trust determinants in computer-mediated transactions: A cross-cultural comparison study," *J. Manag. Inf. Syst.*, 2008, doi: 10.2753/MIS0742-1222240401.
- [3] K. D. Aiken, B. S. Liu, R. D. Mackoy, and G. E. Osland, "Building internet trust: signalling through trustmarks," *Int. J. Internet Mark. Advert.*, 2004, doi: 10.1504/IJIMA.2004.005017.
- [4] A. Belhadi, "Smart Factory Implementation in Moroccan Phosphate Industry," *SSRN Electron. J.*, 2020, doi: 10.2139/ssrn.3637962.
- [5] Undesa, "Managing Knowledge to Build Trust in Government," *Work. Manag. Knowl. to Build Trust Gov.*, 2007.
- [6] T. R. Howe, "A quiet revolution: Developmental psychopathology as a cutting edge framework for understanding abuse and neglect.," *PsycCRITIQUES*, 2016, doi: 10.1037/a0040493.
- [7] J. Rosse, "Digital Is About Speed — But It Takes a Long Time," *MIT Sloan Manag. Rev.*, 2018.
- [8] United Nations, "Managing Knowledge to Build Trust in Government," *Workshop on Managing Knowledge to Build Trust in Government*. 2007.
- [9] D. I. Gertman, R. L. Boring, G. Beitel, and M. Plum, "Characterization process for representation of human factors and human reliability in facility security," in *Transactions of the American Nuclear Society*, 2004.
- [10] K. Ahmed, "Corporate forum brings insight and quality to meet changing trends in providing professional training - Interview with Mr Saim Amin -Manager Training, Corporate Forum," *Pakistan Gulf Econ.*, 2017.

*Editorial Board*

Dr. B.S. Rai,  
Editor in Chief  
M.A English, Ph.D.  
Former Principal  
G.N. Khalsa PG.College,  
Yamunanagar, Haryana, INDIA  
Email: balbirsinghrai@yahoo.ca

Dr. Romesh Chand  
Professor- cum-Principal  
CDL College Of Education,Jagadhri,  
Haryana, INDIA  
Email: cdlcoe2004@gmail.com

Dr. R. K.Sharma  
Professor (Rtd.)  
Public Administration,  
P U Chandigarh, India  
Email: sharma.14400@gmail.com

Dr. Mohinder Singh  
Former Professor & Chairman.  
Department of Public Administration  
K. U. Kurukshetra (Haryana)  
Email: msingh\_kuk@yahoo.co.in

Dr. S.S. Rehal  
Professor & chairman,  
Department of English,  
K.U. Kurukshetra (Haryana)  
Email: srehal63@gmail.com

Dr. Victor Sohmen  
Professor,  
Deptt. of Management and Leadership  
Drexel University Philadelphia,  
Pennsylvania, USA.  
Email: vsohmen@gmail.com

Dr. Anisul M. Islam  
Professor  
Department of Economics  
University of Houston-Downtown,  
Davies College of Business  
Shea Street Building Suite B-489  
One Main Street, Houston,  
TX 77002, USA  
Email: islama@uhd.edu

Dr. Zhanna V.Chevychalova, Kharkiv,  
Associate Professor,  
Department of International Law,  
Yaroslav Mudry National Law University,  
UKRAINE  
Email:zhannachevychalova@gmail.com

Dr. Kapil Khanal  
Associate Professor of Management,  
Shankar Dev Campus,  
Ram Shah Path T.U. Kirtipur, NEPAL.  
Email:kapilkhanal848@gmail.com

Dr. Dalbir Singh  
Associate Professor  
Haryana School of Business, G.J.U.S & T, Hisar,  
Haryana, INDIA  
Email: dalbirhsb@gmail.com

Nadeera Jayathunga  
Senior Lecturer  
Department of Social Sciences,  
Sabaragamuwa University, Belihuloya,  
SRI LANKA  
Email: nadeesara@yahoo.com

Dr. Parupalli Srinivas Rao  
Lecturer in English,  
English Language Centre,  
King Faisal University, Al-Hasa,  
KINGDOM of SAUDI ARABIA  
Email: vasupsr@yahoo.com

## Categories

- Business Management
- Social Science & Humanities
- Education
- Information Technology
- Scientific Fields

## Review Process

Each research paper/article submitted to the journal is subject to the following reviewing process:

1. Each research paper/article will be initially evaluated by the editor to check the quality of the research article for the journal. The editor may make use of iThenticate/Viper software to examine the originality of research articles received.
2. The articles passed through screening at this level will be forwarded to two referees for blind peer review.
3. At this stage, two referees will carefully review the research article, each of whom will make a recommendation to publish the article in its present form/modify/reject.
4. The review process may take one/two months.
5. In case of acceptance of the article, journal reserves the right of making amendments in the final draft of the research paper to suit the journal's standard and requirement.

## Published by

### South Asian Academic Research Journals

A Publication of CDL College of Education, Jagadhri (Haryana)  
(Affiliated to Kurukshetra University, Kurukshetra, India)

Our other publications :

South Asian Journal of Marketing & Management Research (SAJMMR)

ISSN (online) : 2249-877X

SAARJ Journal on Banking & Insurance Research (SJBIR)

ISSN (online) : 2319 – 1422