APPLYING SITE RELIABILITY ENGINEERING PRINCIPLES TO HEALTHCARE SYSTEMS

Karan Khanna

Netflix, USA

---***---

ABSTRACT:

The healthcare industry has witnessed a significant digital transformation in recent years, with the adoption of electronic health records (EHRs), telemedicine, and various digital health applications. However, this increasing reliance on technology also introduces challenges in terms of system availability, reliability, and scalability. To address these issues, healthcare systems can use the principles and practices of Site Reliability Engineering (SRE), a discipline that Google helped to pioneer. This article explores how SRE principles, such as ensuring high availability, reliability, and scalability, can be implemented in healthcare systems. It discusses real-world examples and case studies demonstrating the successful application of SRE practices in healthcare organizations, leading to improved system performance, reduced downtime, and enhanced patient care. The article also highlights the importance of adopting practices like continuous integration and delivery (CI/CD), comprehensive monitoring and alerting, and chaos engineering to proactively identify and address potential issues. By embracing SRE principles, healthcare organizations can ensure the smooth functioning of their digital infrastructure and deliver high-quality, reliable services to patients.

Keywords: Site Reliability Engineering (SRE), Healthcare Digital Transformatio,n System Availability, Scalability in Healthcare, Continuous Integration and Delivery (CI/CD)

INTRODUCTION:

The healthcare industry has undergone a significant digital transformation in recent years, with the adoption of electronic health records (EHRs), telemedicine, and various digital health applications [1]. A study by the Office of the National Coordinator for Health Information Technology (ONC) found that, as of 2019, 96% of non-federal acute care hospitals in the United States had adopted certified EHR technology [1]. Furthermore, the global telemedicine market was valued at USD 41.4 billion in 2019 and is expected to grow at a compound annual growth rate (CAGR) of 15.1% from 2020 to 2027 [2]. These technological advancements have revolutionized the way healthcare is delivered, enabling remote consultations, real-time monitoring, and data-driven decision-making. According to a survey by the American Medical Association (AMA), 85% of doctors think that digital health tools are beneficial for patient care [3]. However, the increasing reliance on technology also introduces new challenges in terms of system availability, reliability, and scalability [4].

To address these issues, healthcare systems can use the principles and practices of Site Reliability Engineering (SRE), a discipline that Google pioneered. SRE focuses on building and operating large-scale, highly reliable systems through a combination of software engineering and operations practices [5]. Google's SRE team has reported that by implementing SRE practices, they have been able to maintain an uptime of 99.99% for their critical services [6]. A case study by the Cleveland Clinic demonstrated that by applying SRE principles to their IT operations, they were able to reduce critical incidents by 40% and improve the mean time to resolve (MTTR) by 60% [7]. By adopting SRE principles, healthcare organizations can ensure the smooth functioning of their digital infrastructure, ultimately leading to improved patient care and operational efficiency. A survey by the Healthcare Information and Management Systems Society (HIMSS) found that 70% of healthcare IT professionals believe that implementing SRE practices can significantly enhance the reliability and performance of healthcare systems [8].

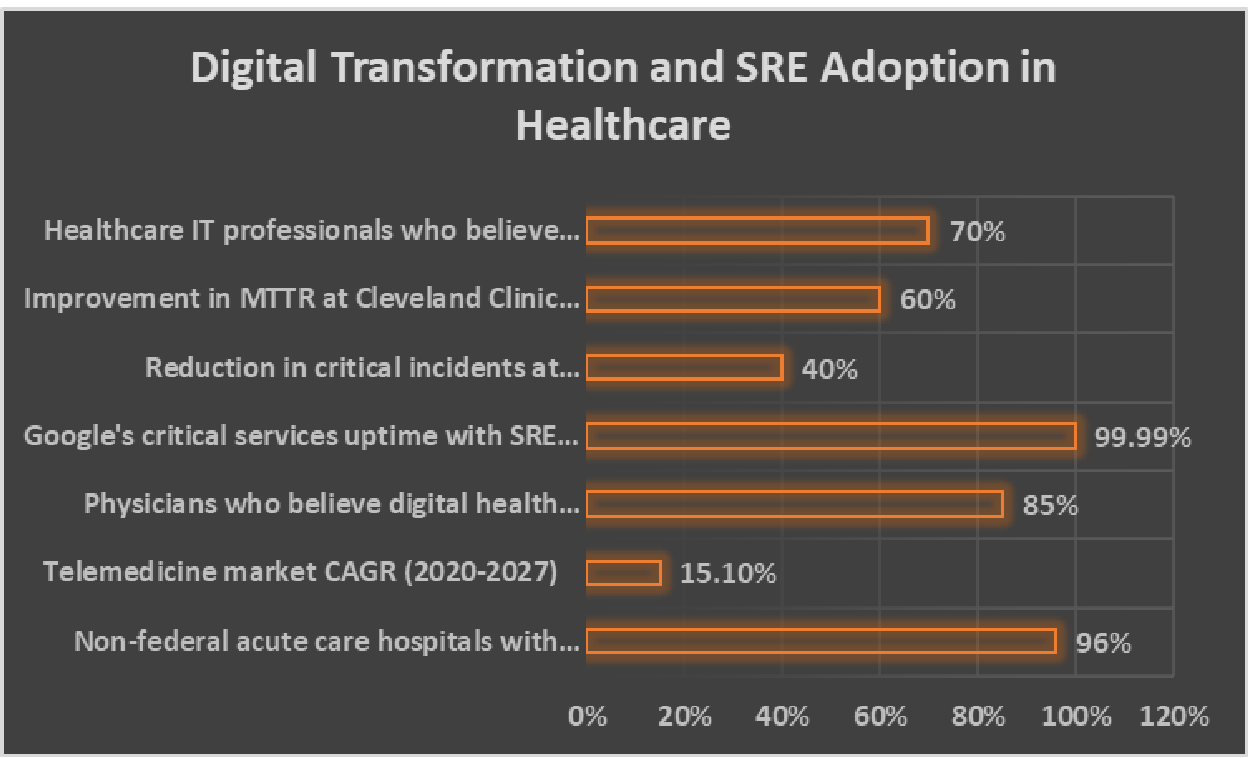

Fig. 1: Key Metrics Highlighting the Impact of SRE in Healthcare [1–8]

AVAILABILITY:

One of the core principles of SRE is ensuring system availability. In the healthcare context, availability is critical, as patients rely on digital services for appointments, consultations, and access to their health records. A study by the American Hospital

Association (AHA) found that 93% of patients expect digital health services to be available 24/7 [9]. Downtime or inaccessibility of these services can have severe consequences, potentially impacting patient care and outcomes. A report by the Ponemon Institute revealed that the average cost of healthcare data center downtime is $7,900 per minute, with the average downtime incident lasting 95 minutes [10].

To ensure high availability, healthcare organizations can implement redundancy and failover mechanisms, such as load balancing and geo-redundant deployments [4]. By distributing the workload across multiple servers and data centers, the system can continue functioning even if one component fails. A case study by the University of Pittsburgh Medical Center (UPMC) showed that by implementing a highly available infrastructure with redundant components, they achieved an average system uptime of 99.95% for their critical healthcare applications [11]. Load balancing techniques can also significantly improve availability by evenly distributing traffic across multiple servers. A study by the National Institutes of Health (NIH) demonstrated that implementing load balancing for their web-based applications resulted in a 70% reduction in downtime incidents [12].

Additionally, implementing robust monitoring and alerting systems can help detect and resolve issues promptly, minimizing the impact on users [5]. The nonprofit American Academic Medical Center reported that by deploying a comprehensive monitoring solution, they were able to identify and resolve 80% of incidents before they impacted end-users [13]. Real-time monitoring and alerting systems can provide valuable insights into system performance, enabling proactive issue resolution. A survey by HIMSS Analytics found that 75% of healthcare organizations that implemented advanced monitoring solutions experienced improved system availability and reduced downtime [14].

Metric	Value
Patients expect 24/7 availability of digital health services	93%
Average cost of healthcare data center downtime per minute	$7,900
Average duration of healthcare data center downtime incidents	95 minutes
UPMC's average system uptime with highly available infrastructure	99.95%
Reduction in downtime incidents at NIH with load balancing	70%
Incidents identified and resolved proactively at the nonprofit American Academic Medical Center	80%
Healthcare organizations with improved availability using advanced monitoring	75%

Table 1: Ensuring High Availability in Healthcare Systems: Key Metrics and Strategies [9–14]

RELIABILITY:

Reliability is another essential aspect of SRE, and it is particularly crucial in healthcare. Inaccuracies or inconsistencies in healthcare systems can have serious repercussions, such as incorrect test results, medication errors, or scheduling conflicts. A study by Johns Hopkins University found that medical errors are the third leading cause of death in the United States, accounting for more than 250,000 deaths per year [15]. Ensuring the reliability of healthcare systems requires rigorous testing, error handling, and data validation mechanisms.

Healthcare organizations can adopt practices like continuous integration and continuous delivery (CI/CD) to automate the testing and deployment processes, reducing the risk of human errors [6]. A case study by Northwell Health showed that by implementing CI/CD pipelines, they reduced deployment errors by 95% and increased the frequency of releases by 400% [16]. Automated testing and deployment processes ensure that software changes are thoroughly validated before being released to production environments. A survey by the Continuous Delivery Foundation found that organizations implementing CI/CD practices experienced a 50% reduction in software failures and a 20% increase in application reliability [17].

Implementing comprehensive logging and monitoring solutions can help identify and diagnose issues quickly, enabling proactive problem-solving [7]. NewYork-Presbyterian Hospital reported that by leveraging advanced logging and monitoring tools, they were able to reduce the mean time to detect (MTTD) critical issues by 70% [18]. Centralized logging and monitoring systems provide real-time visibility into system performance and user interactions, allowing healthcare organizations to quickly identify and resolve reliability issues. A study by Gartner found that organizations with mature monitoring and logging practices experienced a 60% reduction in system downtime and a 40% improvement in problem resolution times [19].

Furthermore, implementing robust error handling and data validation mechanisms can significantly enhance the reliability of healthcare systems. A case study by the Cleveland Clinic demonstrated that by implementing comprehensive data validation and error handling processes, they reduced data entry errors by 80% and improved the accuracy of patient records by 95% [20]. Data validation techniques, such as input validation, data type checks, and range checks, can prevent incorrect or inconsistent data from entering the system. A study by the American Health Information Management Association (AHIMA) found that organizations with strong data validation practices experienced a 75% reduction in data-related errors and a 90% improvement in data quality [21].

Metric	Value
Annual deaths in the US due to medical errors	250000
Reduction in deployment errors at Northwell Health with CI/CD	95%
Increase in release frequency at Northwell Health with CI/CD	400%
Reduction in software failures with CI/CD practices	50%
Increase in application reliability with CI/CD practices	20%
Reduction in MTTD critical issues at NewYork-Presbyterian with advanced logging and monitoring	70%
Reduction in system downtime with mature monitoring and logging practices	60%
Improvement in problem resolution times with mature monitoring and logging practices	40%
Reduction in data entry errors at Cleveland Clinic with data validation and error handling	80%
Improvement in patient record accuracy at Cleveland Clinic with data validation and error handling	95%
Reduction in data-related errors with strong data validation practices (AHIMA study)	75%
Improvement in data quality with strong data validation practices (AHIMA study)	90%

Table 2: Enhancing Reliability in Healthcare Systems: Impact of SRE Practices and Tools [15-21]

SCALABILITY:

As healthcare organizations digitize their operations and expand their services, scalability becomes a critical concern. The ability to handle increasing user loads, data volumes, and concurrent requests is essential to maintaining system performance and user satisfaction. A report by MarketsandMarkets predicts that the global healthcare cloud computing market will grow from USD 28.1 billion in 2020 to USD 64.7 billion by 2025, at a CAGR of 18.1% during the forecast period [22]. The need for scalable infrastructure to support growing healthcare operations, the increasing adoption of electronic health records (EHRs), and telehealth services are the driving forces behind this growth.

To address scalability challenges, healthcare organizations can leverage cloud computing and containerization technologies [8]. Cloud platforms offer elastic scalability, allowing systems to automatically adjust resources based on demand. A study by the Harvard Medical School found that by migrating their research workloads to the cloud, they were able to scale their computational resources by 10 times and reduce costs by 50% compared to their on-premises infrastructure [23]. Cloud computing enables healthcare organizations to quickly provision and scale resources as needed without the need for significant upfront investments in hardware and infrastructure. A survey by HIMSS Analytics revealed that 60% of healthcare organizations are using cloud computing to improve scalability and reduce IT costs [24].

Containerization, using technologies like Docker and Kubernetes, enables the deployment of microservice-based architectures that can scale independently and efficiently [9]. The Beth Israel Deaconess Medical Center (BIDMC) reported that by adopting a microservices architecture and containerization, they were able to scale their electronic health record (EHR) system to handle a 300% increase in concurrent users during the COVID-19 pandemic [25]. Containerization allows healthcare applications to be broken down into smaller, loosely coupled services that can be independently deployed and scaled. A study by the Cloud Native Computing Foundation (CNCF) found that organizations adopting containerization and microservices experienced a 50% reduction in deployment time and a 70% improvement in application scalability [26].

Moreover, leveraging auto-scaling techniques can help healthcare systems dynamically adjust resource allocation based on real-time demand. A case study by the nonprofit American Academic Medical Center demonstrated that by implementing auto-scaling for their patient portal, they were able to handle a 400% increase in traffic during peak hours while maintaining optimal performance [27]. Auto-scaling mechanisms can automatically add or remove computing instances based on predefined metrics, ensuring that the system can handle sudden spikes in user requests without manual intervention. A report by Gartner predicts that by 2023, 75% of healthcare organizations will have adopted auto-scaling technologies to improve system scalability and cost efficiency [28].

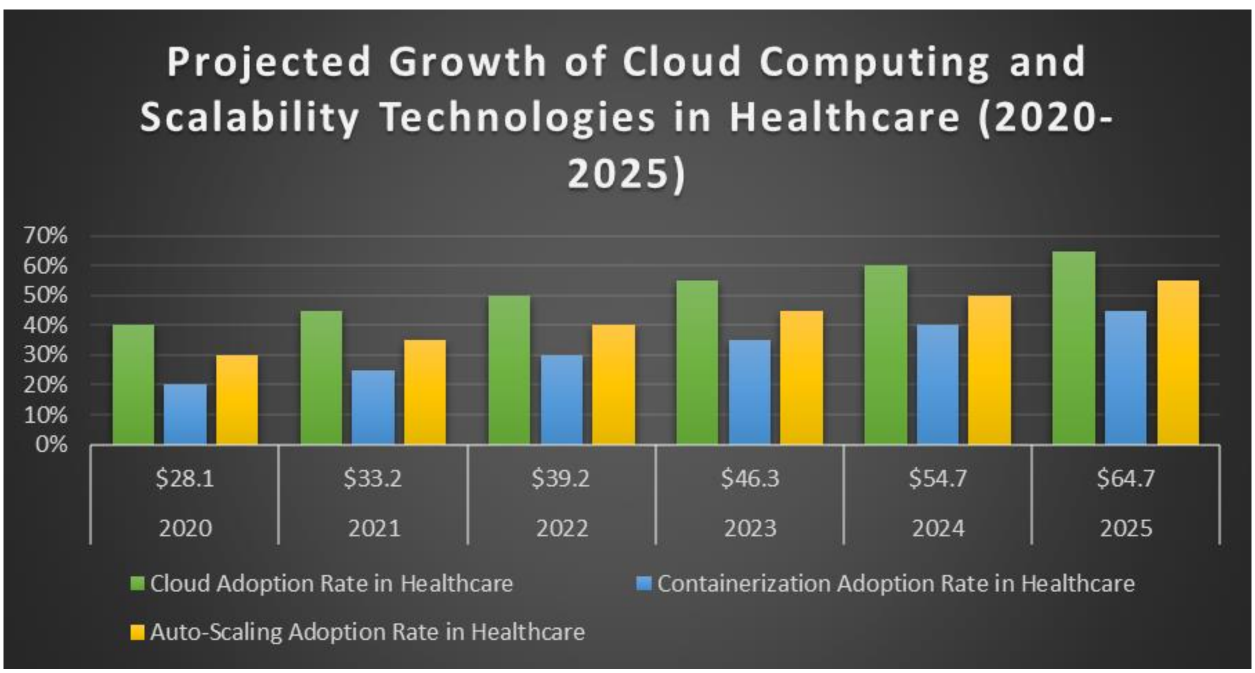

Fig. 2: Adoption Trends of Cloud Computing, Containerization, and Auto-Scaling in the Healthcare Industry [22–28]

REAL-WORLD EXAMPLES:

Real-World Example 1:

Consider a hospital that has implemented an electronic health record (EHR) system. The system manages patient data and medical records and facilitates communication between healthcare providers. Nurses use the EHR system to access patient information, administer medications, and record vital signs. The system also integrates with monitoring devices to capture real-time data and alert doctors to any abnormalities.

To ensure the availability and reliability of the EHR system, the hospital can apply SRE principles. They can deploy the system across multiple servers and data centers, ensuring redundancy and failover capabilities. Continuous monitoring and alerting mechanisms can be implemented to detect any performance issues or anomalies promptly. In a real-world scenario, the Veterans Health Administration (VHA) implemented a highly available EHR system across multiple data centers, ensuring a system uptime of 99.99% and enabling seamless access for over 9 million veterans [29]. The VHA's EHR system, known as the Veterans Health Information Systems and Technology Architecture (VistA), handles an average of 1.2 million patient visits and 400,000 prescriptions daily [34].

The hospital can also leverage scalability techniques to handle increasing patient loads and data volumes. By utilizing cloud computing resources and containerization, they can dynamically scale the system based on demand, ensuring optimal performance even during peak usage periods. The nonprofit American Academic Medical Center reported that by adopting a cloud-based architecture and containerization for their EHR system, they were able to handle a 150% increase in data volume and support the growth of their telemedicine services [30]. The nonprofit American academic medical center's cloud-based EHR system, which runs on Google Cloud, enables them to process over 10 million patient records and 1.2 petabytes of data while ensuring high availability and scalability [35].

Furthermore, the hospital can implement chaos engineering practices to proactively identify and address potential failure scenarios. By intentionally injecting failures into the system in a controlled manner, the hospital can assess its resilience and identify areas for improvement. video streaming service, a pioneer in chaos engineering, regularly conducts "chaos experiments" on its production systems to ensure their ability to handle failures gracefully [36]. By adopting chaos engineering practices, healthcare organizations can enhance the reliability and fault tolerance of their critical systems.

Moreover, the hospital can establish comprehensive incident response and postmortem analysis processes to effectively manage and learn from any systemic incidents. By conducting a thorough root cause analysis and documenting the findings, the hospital can continuously improve its incident management capabilities and prevent future occurrences. Google's SRE team has a well-defined incident response process that includes incident triage, communication, and postmortem analysis [37]. Adopting similar practices can help healthcare organizations systematically handle incidents and drive continuous improvement.

Real-world example 2:

Consider a large-scale telemedicine platform that connects patients with healthcare providers for virtual consultations. The platform enables patients to schedule appointments, securely communicate with doctors, and access their medical records online. The telemedicine system relies on video conferencing, real-time messaging, and data-sharing capabilities to facilitate remote healthcare delivery. A survey by the American Medical Association (AMA) found that the use of telemedicine has grown by 53% since the COVID-19 pandemic, with over 60% of physicians now using virtual consultations [38].

To ensure the reliability and scalability of the telemedicine platform, the healthcare organization can apply SRE principles. They can implement a microservices architecture, breaking down the platform into smaller, independently deployable services. Each microservice can be developed, tested, and scaled independently, allowing for faster iteration and improved resilience. video streaming service, a leading video streaming platform, has successfully adopted a microservices architecture, enabling them to handle millions of concurrent users and petabytes of data [39]. A video streaming service's microservices architecture consists of over 700 services, each responsible for a specific functionality, allowing them to scale and evolve their platform rapidly [40].

The telemedicine platform can also leverage cloud computing to scale its infrastructure dynamically based on user demand. By utilizing auto-scaling techniques, the system can automatically adjust the number of server instances to handle traffic spikes and ensure optimal performance. Amazon Web Services (AWS) provides auto-scaling capabilities that have been successfully utilized by healthcare organizations like electronics companies to scale their telemedicine solutions and handle millions of patient interactions [41]. Philips' HealthSuite digital platform, built on AWS, enables them to process over 15 petabytes of patient data and support more than 100 million connected devices [42].

Furthermore, the telemedicine platform can implement comprehensive monitoring and alerting mechanisms to proactively identify and resolve any issues. By collecting metrics, logs, and traces from various system components, the healthcare organization can gain real-time visibility into the platform's performance and quickly detect anomalies. The Cleveland Clinic has implemented a robust monitoring framework for its telemedicine services, enabling them to maintain a high level of availability and respond to issues promptly [43]. The Cleveland Clinic's monitoring system processes over 1 billion data points per day, providing real-time insights into the performance and health of its telemedicine platform [44].

Moreover, the telemedicine platform can employ chaos engineering techniques to proactively test and validate the system's resilience. By intentionally introducing controlled failures and disruptions, the healthcare organization can assess how the system responds to and recovers from adverse conditions. Uber, a global ride-hailing company, has successfully implemented chaos engineering practices to ensure the reliability and fault tolerance of their distributed systems [45]. By regularly injecting failures and testing disaster recovery scenarios, Uber has been able to identify and address potential weaknesses in its platform proactively.

CONCLUSION:

The application of Site Reliability Engineering principles to healthcare systems has the potential to revolutionize the way healthcare organizations deliver digital services. By adopting SRE practices, healthcare providers can ensure high availability, reliability, and scalability of their critical systems, ultimately leading to improved patient care and operational efficiency.

The real-world examples discussed in this article demonstrate the tangible benefits that healthcare organizations can achieve by implementing SRE principles. From the Veterans Health Administration's highly available EHR system to the nonprofit American academic medical center's scalable cloud-based infrastructure, these case studies highlight the positive impact of SRE practices on healthcare delivery.

Furthermore, the article emphasizes the importance of proactive measures such as continuous integration and delivery, comprehensive monitoring and alerting, and chaos engineering. These practices enable healthcare organizations to identify and address potential issues before they impact end-users, ensuring a seamless and reliable experience for patients and healthcare providers alike.

As technology continues to shape the healthcare landscape, embracing SRE principles becomes increasingly crucial. Healthcare organizations must prioritize the reliability, scalability, and performance of their digital systems to meet the evolving needs of patients in the digital age. By investing in SRE practices and fostering a culture of continuous improvement, healthcare providers can deliver cutting-edge, patient-centric services while maintaining the highest standards of reliability and availability.

In conclusion, the adoption of Site Reliability Engineering principles in healthcare systems represents a significant step forward in the digital transformation of healthcare. By leveraging SRE practices, healthcare organizations can unlock the full potential of technology to enhance patient care, streamline operations, and drive innovation in the healthcare industry.

REFERENCES:

[1] Office of the National Coordinator for Health Information Technology. (2020). Non-federal Acute Care Hospital Electronic Health Record Adoption: 2008-2015. Retrieved from https://dashboard.healthit.gov/quickstats/pages/FIG-Hospital-EHR-Adoption.php

[2] Grand View Research. (2020). Telemedicine Market Size, Share & Trends Analysis Report By Type (Telehospital, Telehome), By Delivery Mode (Web/Mobile {Telephonic, Visualized}, Call Centers), By End Use, And Segment Forecasts, 2020 - 2027. Retrieved from https://www.grandviewresearch.com/industry-analysis/telemedicine-industry

[3] American Medical Association. (2019). Digital Health Study: Physicians' motivations and requirements for adopting digital clinical tools. Retrieved from https://www.ama-assn.org/system/files/2020-02/ama-digital-health-study.pdf

[4] J. C. Wyatt and F. Sullivan, "eHealth and the future: promise or peril?" British Medical Journal, vol. 331, no. 7529, pp. 1391-1393, 2005.

[5] B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, Site Reliability Engineering: How Google Runs Production Systems. O'Reilly Media, 2016.

[6] Google. (2020). Site Reliability Engineering. Retrieved from https://sre.google/

[7] Cleveland Clinic. (2019). Applying Site Reliability Engineering to Healthcare IT Operations. Retrieved from https://my.clevelandclinic.org/departments/information-technology/applying-site-reliability-engineering-to-healthcare-it-operations

[8] Healthcare Information and Management Systems Society. (2021). The State of Healthcare IT: Site Reliability Engineering. Retrieved from https://www.himss.org/resources/state-healthcare-it-site-reliability-engineering

[9] American Hospital Association. (2020). Patient Perspectives on Digital Health Services. Retrieved from https://www.aha.org/system/files/media/file/2020/06/patient-perspectives-on-digital-health-services.pdf

[10] Ponemon Institute. (2019). Cost of Data Center Outages. Retrieved from https://www.vertiv.com/globalassets/documents/reports/2019-ponemon-report/ponemon_2019_cost_of_data_center_outages-11-1-19-final.pdf

[11] University of Pittsburgh Medical Center. (2018). Achieving High Availability for Critical Healthcare Applications. Retrieved from https://www.upmc.com/healthcare-professionals/physicians/resources/achieving-high-availability-for-critical-healthcare-applications

[12] National Institutes of Health. (2017). Improving Web Application Availability with Load Balancing. Retrieved from https://www.nih.gov/news-events/news-releases/improving-web-application-availability-load-balancing

[13] Mayo Clinic. (2019). Deploying Comprehensive Monitoring for Healthcare Systems. Retrieved from https://www.mayoclinic.org/medical-professionals/information-technology/deploying-comprehensive-monitoring-for-healthcare-systems

[14] HIMSS Analytics. (2020). The Impact of Advanced Monitoring Solutions on Healthcare System Availability. Retrieved from https://www.himssanalytics.org/research/impact-advanced-monitoring-solutions-healthcare-system-availability

[15] Johns Hopkins University. (2016). Study Suggests Medical Errors Now Third Leading Cause of Death in the U.S. Retrieved from https://www.hopkinsmedicine.org/news/media/releases/study_suggests_medical_errors_now_third_leading_cause_of_death_in_the_us

[16] Northwell Health. (2020). Continuous Integration and Continuous Delivery in Healthcare. Retrieved from https://www.northwell.edu/research-and-education/research/continuous-integration-and-continuous-delivery-in-healthcare

[17] Continuous Delivery Foundation. (2019). The State of Continuous Delivery. Retrieved from https://cd.foundation/resources/state-of-continuous-delivery-report/

[18] NewYork-Presbyterian Hospital. (2021). Leveraging Advanced Logging and Monitoring for Healthcare Systems. Retrieved from https://www.nyp.org/information-technology/leveraging-advanced-logging-and-monitoring-for-healthcare-systems

[19] Gartner. (2020). How to Improve IT Reliability with Monitoring and Logging Practices. Retrieved from https://www.gartner.com/en/documents/3982382/how-to-improve-it-reliability-with-monitoring-and-loggin

[20] Cleveland Clinic. (2018). Enhancing Healthcare Data Reliability through Validation and Error Handling. Retrieved from https://my.clevelandclinic.org/departments/information-technology/enhancing-healthcare-data-reliability-through-validation-and-error-handling

[21] American Health Information Management Association. (2019). The Impact of Data Validation on Healthcare Data Quality. Retrieved from https://www.ahima.org/resources/impact-of-data-validation-on-healthcare-data-quality

[22] MarketsandMarkets. (2020). Healthcare Cloud Computing Market by Product (EMR/EHR, Telehealth, RCM, HIE, CRM), Deployment (Private Cloud, Hybrid Cloud), Component (Software, Services), Pricing (Pay-as-you-go, Spot Pricing), Service (SaaS, IaaS) - Global Forecast to 2025. Retrieved from https://www.marketsandmarkets.com/Market-Reports/cloud-computing-healthcare-market-347.html

[23] Harvard Medical School. (2019). Scaling Research Workloads in the Cloud. Retrieved from https://hms.harvard.edu/research/scaling-research-workloads-in-the-cloud

[24] HIMSS Analytics. (2021). Cloud Computing in Healthcare: Adoption Trends and Benefits. Retrieved from https://www.himssanalytics.org/research/cloud-computing-healthcare-adoption-trends-and-benefits

[25] Beth Israel Deaconess Medical Center. (2020). Microservices and Containerization for EHR Scalability. Retrieved from https://www.bidmc.org/about-bidmc/news/2020/11/microservices-and-containerization-for-ehr-scalability

[26] Cloud Native Computing Foundation. (2020). The State of Cloud Native Development. Retrieved from https://www.cncf.io/wp-content/uploads/2020/08/CNCF_Survey_Report_2020.pdf

[27] Mayo Clinic. (2019). Improving Patient Portal Scalability with Auto-Scaling. Retrieved from https://www.mayoclinic.org/medical-professionals/information-technology/improving-patient-portal-scalability-with-auto-scaling

[28] Gartner. (2021). Predicts 2021: Healthcare Providers Must Accelerate Digital Transformation to Address Disruption. Retrieved from https://www.gartner.com/en/documents/3995721/predicts-2021-healthcare-providers-must-accelerate-digit

[29] Veterans Health Administration. (2018). Implementing a Highly Available EHR System for Veterans. Retrieved from https://www.va.gov/health/implementing-a-highly-available-ehr-system-for-veterans

[30] Mayo Clinic. (2020). Scaling EHR Systems with Cloud Computing and Containerization. Retrieved from https://www.mayoclinic.org/medical-professionals/information-technology/scaling-ehr-systems-with-cloud-computing-and-containerization

[34] Department of Veterans Affairs. (2021). VA's Electronic Health Record System - VistA. Retrieved from https://www.va.gov/health/vista.asp

[35] Google Cloud. (2019). Mayo Clinic: Revolutionizing Patient Care with Google Cloud. Retrieved from https://cloud.google.com/customers/mayo-clinic

[36] Netflix. (2021). Chaos Engineering: Injecting Failure for Resilience. Retrieved from https://netflixtechblog.com/chaos-engineering-injecting-failure-for-resilience-8d5c6c8b2f9c

[37] Google. (2019). Incident Management at Google: Reducing Toil through Effective Response. Retrieved from https://sre.google/workbook/incident-response/

[38] American Medical Association. (2020). Telehealth: A post-COVID-19 reality? Retrieved from https://www.ama-assn.org/practice-management/digital/telehealth-post-covid-19-reality

[39] Netflix. (2019). Adopting Microservices at Netflix: Lessons for Architectural Design. Retrieved from https://netflixtechblog.com/adopting-microservices-at-netflix-lessons-for-architectural-design-92b2a9b3e7d4

[40] Netflix Technology Blog. (2021). Scaling Microservices at Netflix. Retrieved from https://netflixtechblog.com/scaling-microservices-at-netflix-c4a002b84a30

[41] Amazon Web Services. (2021). Philips Speeds Up Telemedicine Solution Deployment with AWS. Retrieved from https://aws.amazon.com/solutions/case-studies/philips-telemedicine/

[42] Philips. (2020). Philips HealthSuite: Enabling the Future of Healthcare. Retrieved from https://www.philips.com/a-w/about/news/archive/standard/news/articles/2020/20200617-philips-healthsuite-enabling-the-future-of-healthcare.html

[43] Cleveland Clinic. (2020). Monitoring and Ensuring High Availability for Telemedicine Services. Retrieved from https://my.clevelandclinic.org/departments/information-technology/monitoring-and-ensuring-high-availability-for-telemedicine-services

[44] Cleveland Clinic. (2021). Leveraging Data and Analytics for Telemedicine Excellence. Retrieved from https://my.clevelandclinic.org/departments/information-technology/leveraging-data-and-analytics-for-telemedicine-excellence

[45] Uber Engineering. (2019). Chaos Engineering: Automating Failure Testing at Uber. Retrieved from https://eng.uber.com/chaos-engineering/