October 2022

White Paper

The Business Case for Observability

Rationale, Requirements, Rewards, and Recommendation

Abstract

IT environments with complex, modern architectures are constantly changing and expose us to potential system failures in unpredictable ways. Tracing an event to its origin in these distributed systems requires massive amounts of data and thousands of processes running on the cloud, on-premises, or both. Conventional monitoring techniques and tools cannot provide contextual tracing to assist IT teams in gaining actionable insights leading to visibility gaps and performance challenges.

With growing data-driven applications and multi-cloud infrastructure, organizations need to have not only a broad visibility of its data landscape but also action-oriented insights. Organizations need to observe infrastructure (where data is?), application (what components are using this data?), usage (in what context?), security (how liable?), etc. The objective is to proactively identify, control, prevent, and remediate data outages by quickly using a unified observability platform.

The need for observability is driven by multiple forces as listed below.

1. Organizations are rapidly adopting **distributed architecture and software development models** bringing more products and services to the market faster than ever. They include cloud-native – microservices, Docker containers, Kubernetes, and serverless functions, hybrid/multi-clouds, agile development, continuous integration and continuous deployment (CI/CD), DevOps, DevSecOps, multiple programming languages, mobile edge, Internet of Things (IoT), and AIOps changes

2. The need for **enhanced customer experience, innovation, business agility, and quality** as organizations continue to embrace digital transformation initiatives

3. **Operational challenges** with remote cross-functional teams have **contributed to significantly longer triage times** impacting customer service and revenue

Adoption of an observability platform is imperative for organizations to achieve end-to-end visibility of their technology stacks to ensure a highly performant, resilient, and self-healing system thereby avoiding expensive disruptions to business.

The observability business benefits include:

- **Reduced IT Costs:** Understand complex systems at all levels and leverage a unified platform to reduce the costs of monitoring cloud infrastructure with significantly faster root cause analysis.

- **Increased IT Uptime:** Proactively identify performance issues before they cause a severe outage and lead to a loss of revenue.
- **Accelerated Time to Market:** Link together business and technical data to stay competitive in a fast-moving market while maintaining a competitive edge by increasing visibility into deployments and their effect on application performance.
- **Increased customer conversion:** Ship products and services to customers faster to drive higher customer satisfaction by delivering responsive applications with a seamless user experience.
- **Reduced security risks:** Leverage existing observability data for a complete view of application threats and vulnerabilities.

Many topical observability business cases include cloud migration, DevOps, DevSecOps, shift-left testing, customer/employee/digital experience monitoring, security analytics, compliance, hybrid cloud monitoring, machine learning, real-time business intelligence, and log analysis & correlation.

Datadog is recognized as an observability leader in Gartner's 2022 Magic Quadrant[11] and rated as "exceptional" in focus and execution by GigaOm Radar for Cloud Observability Solutions 2022[12]. Datadog supports OpenTelemetry, offers a holistic observability platform with 500+ out-of-the-box integrations, and built-in security monitoring capabilities.

Problem Statement

On average, organizations face a $35.5 million annual revenue loss due to delays in application release times[1].

Organizations need a high-performant, efficient, flexible, and fault-tolerant technology landscape. To achieve this, the infrastructure and application tech stack must be "observable"— i.e., ability to proactively detect and resolve issues before it impacts the customer; at a minimum, swift identification of the root-cause. Current monitoring tools provide metrics, logs, and traces for IT teams. Given the distributed architecture, it would be nearly impossible to identify the root cause at the speed of business needs without observability.

As the technology landscape continues to increase in complexity, it will be imperative for organizations to implement observability platforms to accelerate time-to-market, reduce IT costs, increase uptime, and improve security posture.

The top two imminent drivers for observability implementation include:

1. **Enhance Customer Experience**
    - Customer experience is a boardroom agenda. Customers expect "always on" applications with little latency. These applications can includeseveral microservices talking to each other behind the scenes, making them immensely complex from an operations perspective.
    - Customer loyalty is based on quality of service and not on the number of years with an organization. Nearly 80% of companies are putting their customer loyalty at risk[2]. Applications are accessed through web interfaces or a mobile app that provides rich and interactive applications. Here are a couple of relevant cases:
        - 96% of customers agree that customer service plays a vital role in their choice of loyalty to a brand[3].
        - 9x revenue is generated with a one star increase in an app store[4]
        - A site that loads in 1 second has a conversion rate 3x higher than a site that loads in 5 seconds[5].
        - Around 50% of customers say they would switch to a new brand after a bad customer service experience[6].

2. Mitigate outage and brand erosion

- Application and service outages can cause significant financial loss and reputational damage. Here are some examples:
    - Organizations are losing $5.1 million annually due to a lack of balance between cost/resources and performance of managing digital services[7].
    - As stated in the 2022 state of observability report[8], organizations with observability reduce costs associated with downtime by 90%, from $23.8 million for beginners to $2.5 million for leaders.
    - Per DEJ report[9], organizations with observability see a 75% reduction in Mean Time to Resolution per incident.

Determining the business context is crucial to achieving true observability. Organizations need to identify specific business outcomes (KPIs) for observability implementation such as:

- **Financial** (e.g., increase revenue from new digital services and new product development, accelerate time-to-market, increase customer satisfaction, reduce customer churn, increase customer conversion, reduce security risks, increase productivity, and reduce cost through automation)
- **Technology** (e.g., create actionable insights at scale, proactive detection, reduced MTTR, quick resolution, increased application releases, ensure optimal customer experience)
- **Business Process** (e.g., understand business context of IT performance, enable innovation, establish predictable operations, manage change and complexity)
- **People** (e.g., reduce engineer toil, alignment of internal service level objectives [SLOs] and internal service level indicators [SLIs] with the business metrics)

Background

Organizations are under immense competitive pressure to innovate their products and services at the speed of business and/or customer need. This requires them to adopt modern development practices and distributed architectures like DevSecOps, containers, serverless functions, cloud, etc. Traditional monitoring will not keep up with high-frequency deployments of application components in diverse languages and for a varying period of time (e.g., sub-seconds for serverless functions).

A brief history of how the technology landscape has transpired over the past two decades is necessary to not only understand the limitations of the traditional monitoring tools, but also support the need for a holistic observability platform.

Application

The application architecture has evolved from *monolithic, multi-tier architecture* to *service-oriented architecture (SOA)* to *microservices.*

With a multi-tier architecture, separate physical layers — presentation, business logic, and database — allowed IT teams to modify and update the layers independently (e.g., change Oracle to SQL Server at the database layer). The presentation layer had user interface elements and presentation applications. The business logic layer included the logic and all parsing data requirements between presentation and database layers.

The SOA, more granular than the multi-tier architecture, was driven by standardized service contracts and registry; it provided autonomy and abstraction.

With the growing pains of the monolithic architectures (e.g., single technology stack with obsolete libraries, long code build and testing times for code deployment) came the lightweight version of the service-oriented architecture called microservices (fine-grained SOA at web scale) enriched with other recent architectures like:

- Event-driven architecture for event stream processing (e.g., IoT device on an oil tank in a refinery would inform the change in temperature)
- Domain-driven design (DDD) providing the bounded context with a single responsibility for strict consistency

Salient benefits of microservices architecture include:

- Quick feature releases as you can test, build, and deploy frequently and quickly given the very small code base
- Scale quickly and independently as needed

- Freedom to use any type of language or database or web application server
- Flexibility in integrating microservices translates to enterprise's flexibility

Infrastructure

On the infrastructure front, organizations moved away from buying/leasing physical machines to virtual machines (one of the first public cloud offerings) through cloud services with Cloud Service Provider (CSP) provided cost savings (shifting capital expenditure to operational expenditure), security, multi-region autoscaling, and load balancing if needed. Applications were deployed on a host operating system using configuration tools like Chef or Puppet. To deploy faster than VMs, immutable containers (lightweight standalone images) are used to hold the code and dependencies to ensure full operating system level virtualization. An orchestration engine (e.g., Kubernetes) is used to manage containers (e.g., can scale up or down and be made fault-tolerant). Abstracting all the infrastructure configuration management around a service, serverless computing (CSP fully manages the resources) focuses on business logic with a very short lifetime. VMs, containers, and serverless form the basic building blocks behind any application or service hosted in the public cloud.

Database Management

On the database front, traditional monolithic architecture caters to "structured" transactional data. Hence, most databases were relational database management systems (RDBMS). With the ever-growing data from diverse sources (e.g., social media, sensors), it is much more than relational data. For example, shopping carts and product catalogs need high throughput and latency with key-value database, fraud detection and recommendation engine for graphical navigation between data with graph database, content management and personalization need document storage with document database, etc.

Today's environments are only growing more dynamic and complex, involving a range of standards, platforms, and technologies comprising of millions of components. In addition, application developers now working remotely making online collaboration a strategic imperative. Although there is an increase in complexity with microservices deployment (11.4x increase in complexity after deploying microservices[1]), organizations have seen an 8.7x increase in number of new software releases and updates[1] – per a recent DEJ study[1].

Traditional monolithic architecture has been replaced with distributed architecture. Microservices subsumed a three-tier monolithic application tier; large-scale applications are rearchitected with hundreds of microservices that "talk" to one another. Each service can run on its own container. Thus, cloud and container technologies have reshaped the underlying infrastructure and software development and operations have become more dynamic. Today, organizations leverage hybrid and multi-cloud environments by using distributed architectures with microservices orchestration platforms that implement continuous integration and continuous delivery (CI/CD) DevSecOps processes to anticipate and adapt to ever-changing customer expectations. These dynamic IT operation environments expose organizations to unpredictable system failures or security breaches (e.g., one incorrect open-source version in a microservice would create a cascading effect; also, in a recent 2022 survey[3], 59% said security is their biggest concern regarding continued use of Kubernetes and containers[10]); these challenges cannot be handled by traditional monitoring strategies and tools, leading to visibility gaps and performance challenges.

Traditional monitoring is usually applied to monolithic IT environments where a defined set of applications, infrastructure, integrations, and input/output workflows have defined expectations on what to measure and when. In traditional monitoring, IT teams track a system's health and performance using metrics, logs, and traces based on known system parameters (e.g., CPU). However, distributed applications featuring ephemeral infrastructure (e.g., shifting Kubernetes clusters, serverless services) are replete with possibilities where hidden data dependencies can surface perplexing IT teams and disrupting business. Observability offers a way to track "unknown" or "in the dark" phenomena (no visibility) enabling end-to-end visibility of the modern technology stack. It focuses on "what" by providing visibility that makes it clear when performance issues or bottlenecks occur. With the distributed architecture, we need to look for "unknown" unknowns (e.g., shifting and updating Kubernetes clusters in the cloud or a growing list of APIs or serverless services where nothing is constant). IT teams cannot analyze or innovate what they cannot see.

Recommendation

It is highly recommended for organizations to consider implementing an observability platform if one is not already in place. Datadog is the leader in the observability space and is in Gartner's Leadership Quadrant[11] in the Magic Quadrant for application performance monitoring (APM) and observability. Per GigaOm's 2022 Report[12], Datadog is one of the two observability vendors to receive an "exceptional" rating: outstanding focus and execution for both vendor positioning and evaluation metrics comparison charts.

A brief explanation of the foundational capabilities required to achieve enterprise observability at scale are detailed below along with a description of Datadog's features supporting them.

1. **Complete Coverage and Deep Visibility of Technology Stack:** Given the dynamic technology landscape, organizations need to have complete and continuous visibility into the health and performance of all layers of their technological environment. Irrespective of the environment — be it on-premises, hybrid, multi-cloud, or IoT — organizations should be able to deploy observability components anywhere.

    - Datadog offers deep visibility into your technology stack by providing a plethora of out-of-the-box metrics to evaluate your infrastructure and individual host using host maps. Flame graphs help developers identify service dependencies to analyze individual application traces to fix bottlenecks in applications. As a leader in the observability space, it provides complete visibility into infrastructure performance with effortless deployment, minimal maintenance, and wide coverage of environments (on-premises, hybrid, IoT, and multi-cloud environments).

2. **Unified observability platform:** A unified monitoring and observability platform is defined as implementing consistent monitoring processes, workflows, and standards across the organization. Teams employ a centralized platform on which to collect, analyze, alert, and graph their data to gain a comprehensive view of the health and performance of the systems that underpin the business. This democratizes data so that anyone can immediately access data any time and use it in a way that is correlated to the other parts of your business. This eliminates the time-consuming barriers associated with legacy monitoring tools.

    - Datadog helps organizations visualize metrics, traces, logs, and more in a single platform.
        - Metrics without Limits™, enables you to ingest all your metrics, decide which metrics to index that are valuable for your query, and update any time.

- Network Performance Monitoring (NPM)overview dashboard allows you to obtain a unified, high-level view of key network health and performance across different facets of your distributed network.
- Correlate metrics, traces, logs, and network data to get a unified view into your container-based microservices.
- Visualize the health of your serverless architecture—including APIs, queues, databases, and more—in a single unified platform. With AWS SAM and CDK integrations, Datadog helps organizations get deep visibility into your serverless environment.
- Correlate query metrics with database and infrastructure metrics. View query metrics alongside system performance metrics from your Postgres, MySQL and SQL Server databases.
- Leverage out-of-the-box integration with the rest of the Datadog platform including dashboards, monitors, SLO tracking, and advanced formulas and functions.
- Incorporate CI Visibility into your existing monitoring workflow.
- Leverage turn-key integrations with CI providers such as GitLab, Jenkins, CircleCI, etc.
- Automatically alert your teams about any problems in your pipelines.
- Collaborate on investigations into pipeline problems and test issues with Datadog Notebooks and incident management.
- Streamline root-cause investigations via a unified observability tool that leverages built-in integrations for Slack, PagerDuty, and Source Code.
- Maintain efficient DevOps practices while implementing robust threat detection and incident response workflows.
- Enable developers to focus on building new features—not securing your environment.
- Leverage the existing Agent and integrations to achieve new security visibility.
- Utilize unified user journey, triage, and resolve issues with an end-to-end view of user experience.
- Simulate user journeys with Datadog Browser Tests and correlate results with real-time user experiences.

  - Pivot from RUM data to request traces and logs for complete context when investigating issues.
  - Unify full-stack monitoring in a single platform for front-end and back-end development teams.

3. **OpenTelemetry Application Management (APM) Support:** OpenTelemetry is a collection of tools, APIs, and SDKs that organizations use to instrument, generate, collect, and export telemetry data (metrics, logs, and traces). Openness drives innovation and choice, making it a good strategy for an organization's observability journey. OpenTelemetry works towards the common goal of developing an industry wide standard for software instrumentation. Modern cloud-native applications and data are distributed which can make it difficult to compile the data you need into a single database. OpenTelemetry solves this problem by tracing and extracting data cross-platform.

    - OpenTelemetry will enable any company—with any stack, any infrastructure platform, and any monitoring provider—to gather observability data from all systems including distributed traces, metrics, and logs. Because OpenTelemetry is vendor-neutral, companies will be able to migrate their observability data between monitoring backends more easily, without vendor lock-in.
    - Organizations partner with the OpenTelemetry community to build the foundation for auto-instrumentation of applications across languages and frameworks. Datadog provides two paths to move tracing data to Datadog's back-end using OpenTelemetry collector with Data exporter or Datadog Agent. The goal of auto-instrumentation is to make it possible to collect comprehensive telemetry data from applications without making manual changes to your code.

4. **Broad Support for out-of-box integration with other technologies:** Integration is critical as you need to bring together all the metrics and logs from your infrastructure to a unified observability platform to gain insight. This enables organizations the ability to provide options to zoom in (look at specific service/app) or zoom out (look at the holistic tech stack). Integrations come in a variety of forms such as an agent-based, crawler-based (e.g., Slack, AWS, PagerDuty), and API-driven library integrations (e.g., Python). It is ideal to have custom integrations allowing in-house systems to send data.

    - In addition to the OpenTelemetry partnership, Datadog provides 500+ built-in integrations for full visibility into your network, endpoints, SaaS applications, and endpoints. It easily deploys and scales with Docker and Kubernetes integrations and supports extensible, out-of-the-box security integrations. It also enables collaboration with multiple teams via integrations with ticketing portals, chat systems, and remediation tools. Organizations can include alerts into their existing workflow through integrations with notification tools like PagerDuty and

Slack. From a DevOps perspective, IT teams can link error stack traces to relevant source code through integrations with Github, Gitlab, and Bitbucket.

5. **Real User Monitoring (RUM):** 47% of consumers expect websites to load in two seconds or less and 40% will abandon a page that takes three or more seconds[13]. With a customer's low tolerance for unsatisfactory digital experiences, organizations need an observability platform that effectively observes and understands how real users interact with a digital interface (web application or mobile application) and whether their experience is satisfactory.

    - A positive customer experience encounter can increase customer spending by up to 140%[14]. DevOps and SRE teams use Datadog RUM for APM to obtain visibility into every step of the customer's journey to identify front-end issues and how to resolve those errors. Every user's journey is correlated seamlessly with synthetic tests, backend metrics, traces, logs, and network performance data. This automatic correlation of the front-end and back-end enables quick resolution with minimal escalation.
        - Understand user experiences with critical performance data such as Core Web Vitals to visualize and identify the root cause of slow load times, whether it's an issue with the code, the network, or the infrastructure.
        - Resolve JavaScript errors by pinpointing the root cause of an issue down to the line of code.
        - Group high-volume application errors into a small number of issues with error tracking.
        - Identify customer segments using tags and automatically collect and correlate every resource, error, and user action during a user's session.
        - Analyze usage across your mobile and web applications by browser, country, or any custom dimension with RUM Analytics.
        - Contextualize user sessions and capture business-critical user actions such as checkout button clicks.
        - Ingest custom metrics that are specific to your business and correlate them with real-time product analytics.
        - Triage and resolve issues with an end-to-end view of user experiences.
        - Simulate user journeys with Datadog Browser Tests and correlate results with real-time user experiences.
        - Pivot from RUM data to request traces and logs for complete context when investigating issues.
        - Unify full-stack monitoring in a single platform for front-end and back-end development teams.

6. **Synthetic Monitoring:** While RUM tracks customer satisfaction in real time, synthetic monitoring runs predefined scripts to proactively track website performance including service availability, benchmark performance against competitors, and response times. This allows to quickly detect user-facing issues with API and browser tests to proactively enhance user experience. This is useful for service level objective (SLO) reporting and audit.

    - Adding Datadog synthetic browser tests to CI pipelines proactively monitors endpoints by automating browser testing and continuous monitoring for errors in applications. This helps developers to deploy code with confidence. Multistep API tests verify end-to-end workflows and also validate all layers (e.g., HTTP, SSL, DNS, WebSocket, TCP, UDP, ICMP, and gRPC) from several worldwide locations. It allows capturing critical transactions with no code, by using a fully hosted web recorder that helps monitoring key workflows with step-by-step screenshots. Browser tests intelligently detect UI changes and update automatically to reduce alert fatigue.
        - Incorporate Synthetic tests in your CI pipelines for early issue detection and remediation.
        - Evaluate the state of your production environment after each deployment to identify regressions and automate rollbacks.
        - Eliminate false positives with intelligent self-maintaining tests.
        - Run tests from locations around the world.
        - Test all stages of development and throughout your CI/CD pipelines.
        - Minimize downtime and streamline collaboration by eliminating the need for disparate teams to maintain separate testing scenarios.
        - Datadog Synthetic Monitoring enables application engineering, quality assurance, and DevOps/SRE teams to launch new features in production with confidence and increased release velocity while maintaining software quality.

7. **Unlimited data retention and historical analysis:** Amazon's CTO has stated[15], "failures are a given and everything will eventually fail over time: from routers to hard disks, from operating systems to memory units." When failure happens, organizations need a robust process for conducting root cause analysis to identify and alter processes to ensure failures are not repeated. Granular and precise data (not aggregate, for example) from prior history is needed to investigate or correlate with other issues, even on infrastructure that does not exist anymore. Therefore, a "time travel" function to retrieve historical data would be useful to help mitigate future risks.

- Organizations today need to manage petabytes of logs each day from a myriad of systems, applications, and cloud services. This large volume makes retaining all the logs for analysis cost prohibitive. Business seasonality and incidents (e.g., unplanned outages) lead to large variations in log volume and cost. An effective log management solution needs to address these large variations in volume and value to solve the associated pain points. Traditional options would be to either avoid cost by decreased log volume or accept excessive cost storing all logs. Neither of these options are satisfactory. Datadog decouples ingestion and indexing costs. Logging without Limits, a Datadog feature, ingests all logs without filtering, parses all logs at the ingestion time, provides live tail ingestion of logs for real-time visibility during an outage or deployment, and maintains full history for audit or historical analysis.

8. **Security Monitoring:** With complex, multi-cloud environments, organizations see their attack surfaces exponentially expanding. Data security threats include social engineering, malicious insiders, distributed denial of service (DDoS), unpatched vulnerabilities, compromised and weak credentials, misconfigured cloud services, vulnerable mobile devices, etc. It is difficult to effectively secure an environment with limited security data. Organizations need to stream real-time data to enable secure visibility to prevent threats, data loss, fines, and lawsuits to protect their reputation.

    - Datadog's Cloud Security Platform delivers real-time threat detection and continuous configuration audits across the entire production environment, thereby bringing speed, scale, and security to an organization. The Cloud Security Platform is built on top of Datadog's observability platform, which breaks down silos between security and DevOps teams and aligns them to shared organizational goals.

        - Datadog supports real-time detection by remaining vigilant with "always-on" security monitoring that detects attacks at any time, not on a schedule. It leverages continuous scanning to identify misconfigurations, as well as suspicious file and process activity, in real time.

        - Datadog provides the ability to view and correlate context-rich security signals to reconstruct attack paths and secure environments. In addition, Datadog helps assess the business impact of security incidents by pivoting seamlessly between security signals and full-stack telemetry.

        - Datadog's Cloud Security Platform derives real time security insights from detailed observability data logs, metrics, and traces that help DevOps and security teams collaborate to secure their production environment by leveraging the same integrations and lightweight agents that the DevOps teams already use. It analyzes your events for attacks, monitor suspicious file and process activity, and scans the cloud resources for misconfiguration using built-in and custom rules. Datadog creates a security signal for each misconfiguration attack or a threat intelligence match which the platform correlates to the most important signals so teams can quickly reconstruct attack data.

- Datadog's Cloud SIEM(Security Information and Event Management), which is a part of the Datadog Cloud Security Platform, provides robust threat detection for dynamic, cloud-scale environments. Cloud SIEM helps analyze operational and security logs in real time—regardless of their volume—while utilizing curated, out-of-the-box integrations and rules to detect threats. Developers, security, and operations teams can also leverage detailed observability data to accelerate security investigations in a single, unified platform.

Conclusion

With fluctuating customer expectations and the growing complexity of the technology landscape, it is clear that traditional monitoring tools are not sufficient. Observability can address these issues while helping organizations achieve their business objectives.

Observability drives more reliable development with faster issue resolutions and less downtime. An observability platform will:

- Provide knowledge and confidence to DevOps teams to quickly deliver solutions to meet today's business needs.
- Improve developers' productivity, operations' stability, DevOps' visibility, and business stakeholders' optimization goals.
- Improve UX, shortens mean time to repair (MTTR), improves service availability, detects anomalous application behavior, and forecasts future problems.
- Provide application performance through transparency, visibility, and context to improve business goal delivery.
- Reduces infrastructure and developer toil with automation.

Datadog is the leader in the observability space and is listed in Gartner's Magic Quadrant[11] for application performance monitoring (APM) and observability. As referenced in GigaOm's 2022 report[12], Datadog is one of two observability vendors to be rated as exceptional: outstanding focus and execution for both vendor positioning and evaluation metrics comparison charts. Datadog is a major contributor of OpenTelemetry, a vendor-neutral standard with a common goal of developing an industry wide standard for software instrumentation, along with its out-of-the-box 500+ integrations that comes with a host of tagged metrics. Since cybersecurity is a boardroom agenda, with its built-in security monitoring capabilities, Datadog can send observational data to its cloud SIEM product.

Datadog offers a holistic observability platform that addresses organizations' urgency of adopting an observability solution to deliver pressing business outcomes and to stay competitive with the ever-changing, complex technology stack. By implementing an observability platform, your organization can accelerate time-to-market, reduce IT costs, increase uptime, and improve security posture.

To start implementing Datadog's most popular capabilities or to learn about best practices to gain maximum utility from the platform, read our Datadog Best Practice's blog.

References

1. Simic, Bojan. "Observability done right— Practices of leading organizations.", *Digital Enterprise Journal*, 4 Mar. 2022, www.dej.cognanta.com/2022/03/04/observability-done-right-practices-of-leading-organizations/.
2. "Cisco AppDynamics Brandvoice: New Harvard Business Review Study Reveals Nearly 80% of Companies Are Putting Customer Loyalty at Risk." *Forbes*, Forbes Magazine, 9 June 2021, https://www.forbes.com/sites/appdynamics/2019/06/04/new-harvard-business-review-study-reveals-nearly-80-of-companies-are-putting-customer-loyalty-at-risk/?sh=190e62ab20fb.
3. Morgan, Blake. "50 Stats That Prove The Value Of Customer Experience.", *Forbes*, 24 Sept. 2019, www.forbes.com/sites/blakemorgan/2019/09/24/50-stats-that-prove-the-value-of-customer-experience/?sh=56a915574ef2.
4. "Mobile App Ratings and Reviews: Where to Start and How to Win.", *Apptentive*, 2020, go.apptentive.com/rs/170-TZF-108/images/Ratings-Reviews-1.pdf.
5. Weigand, Michael. "Site Speed is (Still) Impacting Your Conversion Rate.", *Portent*, 20 Apr. 2022, www.portent.com/blog/analytics/research-site-speed-hurting-everyones-revenue.htm.
6. Bishop, Court. "31 customer service statistics you need to know." *Zendesk*, 28 Mar. 2022, https://www.zendesk.com/blog/customer-service-statistics/#:~:text=50%25%20of%20consumers%20will%20switch,and%20take%20their%20business%20elsewhere.
7. Simic, Bojan. "10 Key Takeaways From DEJ's 2022 Observability Market Study", *Digital Enterprise Journal*, 17 May 2022, https://www.dej.cognanta.com/2022/05/17/10-key-takeaways-from-dejs-2022-observability-market-study/.
8. "The State of Observability 2022.", *Splunk*, 2022, www.splunk.com/en_us/form/state-of-observability.html.
9. Simic, Bojan. "Identifying Top Vendors for Enabling Observability.", *Digital Enterprise Journal*, 24 May 2022, www.dej.cognanta.com/2022/05/24/identifying-top-vendors-for-enabling-observability/.
10. "Kubernetes Security: 8 Best Practices to Secure Your Cluster", *Tigera*, www.tigera.io/learn/guides/kubernetes-security/.
11. "Datadog Named a Leader in the 2022 Gartner® Magic Quadrant™ for Application Performance Monitoring and Observability", *CIVISION PR Newswire*, 10 Jun. 2022, https://www.prnewswire.com/news-releases/datadog-named-a-leader-in-the-2022-gartner-magic-quadrant-for-application-performance-monitoring-and-observability-301565941.html.
12. Clarke, Sue. and Williams, Ron. "GigaOm Radar for Cloud Observability Solutions v2.0.", *GigaOm*, 25 Mar. 2022, research.gigaom.com/report/gigaom-radar-for-cloud-observability-solutions/.
13. Patel, Neil. "How to Increase Page Speed.", https://neilpatel.com/blog/speed-is-a-killer/.
14. McLain, Simon. "Creating effective customer experience strategies.", *Deloitte*, www2.deloitte.com/us/en/pages/operations/solutions/creating-effective-customer-experience-strategies.html.
15. "Reliability Pillar: AWS Well-Architected Framework", *AWS*, docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/failure-management.html.