



IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 12 Issue: IX Month of publication: September 2024 DOI: https://doi.org/10.22214/ijraset.2024.64151

www.ijraset.com

Call: 🕥 08813907089 🔰 E-mail ID: ijraset@gmail.com



# Designing Tomorrow's Observability: A Software Architect's Guide to Building Effective Monitoring Solutions

Madhu Garimilla Broadcom, USA



Abstract: This comprehensive article explores the evolving landscape of observability in software development, focusing on strategies for designing effective monitoring solutions. It delves into the three pillars of observability - logs, metrics, and traces - and their implementation in modern distributed systems. The paper discusses key performance indicators, tool selection, data collection methods, and best practices for implementing observability. Drawing on industry surveys, case studies, and research reports, it highlights the importance of aligning observability with business goals, automating processes, and continuously improving strategies. The article provides practical insights for software architects and engineers on building robust, scalable observability frameworks that can significantly enhance system reliability and performance.

Keywords: Observability, Microservices, Distributed Tracing, Performance Monitoring, DevOps

# I. INTRODUCTION

Observability has emerged as a critical aspect of system architecture in the ever-evolving software development landscape. Effective monitoring solutions ensure system reliability and provide deep insights into application performance. This article delves into the principles of designing observability frameworks, focusing on best practices, tools, and techniques that software architects can leverage to build robust monitoring solutions.

Observability has evolved significantly over time. Initially, simple logging and basic performance metrics sufficed. The need for more sophisticated observability became evident as complex systems grew, especially with the advent of microservices and distributed systems.



ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538 Volume 12 Issue IX Sep 2024- Available at www.ijraset.com

Today, observability encompasses comprehensive logging, detailed metrics, and intricate tracing. Looking to the future, observability will increasingly integrate AI and machine learning for predictive analytics, self-healing systems, and more proactive monitoring approaches.

Recent industry trends emphasize the significance of observability. According to a 2023 survey by the Cloud Native Computing Foundation (CNCF), 76% of organizations reported that improved observability led to a significant reduction in mean time to resolution (MTTR) for production issues [1]. Furthermore, the same study predicts that by 2025, 70% of organizations will have implemented AI-enhanced observability solutions, up from less than 15% in 2021 [1].

Modern observability frameworks incorporate a wide array of tools and technologies. For instance, the Elastic Stack (formerly ELK Stack) is widely used for log management and analysis. According to the New Relic 2023 Observability Forecast, 88% of organizations are in some stage of observability implementation, with 31% having achieved full-stack observability [2].

As we look to the future, the field of observability is poised for further innovation. Emerging trends include [2]:

- 1) AIOps: AI-driven operations for predictive maintenance and automated issue resolution
- 2) Continuous Verification: Real-time validation of system behavior against expected outcomes
- 3) Observability as Code: Defining and managing observability configurations through code
- 4) OpenTelemetry Adoption: Standardized, vendor-neutral telemetry collection and instrumentation

By embracing these advanced observability practices, organizations can significantly improve their system reliability and performance. The New Relic study found that organizations with mature observability practices are 2.9 times more likely to identify issues before they impact customers and 4.5 times more likely to accelerate their speed to market for digital applications and services [2].

Year	Observability Stage	Description	Adoption Rate (%)
1990s-2000s	Basic Logging	Simple performance metrics	100
Early 2010s	APM Tools	Application Performance Management	80
Mid 2010s	Distributed Tracing	Projects like Zipkin and Jaeger	60
Late 2010s	Three Pillars	Logs, metrics, and traces	40
Early 2020s	AI/ML Integration	Advanced analytics and root cause analysis	15
2025 (Projected)	AI-Enhanced Solutions	Predictive analytics and self- healing systems	70

## Table 1: Evolution of Observability Practices: Adoption Rates Over Time [1, 2]

#### II. UNDERSTANDING OBSERVABILITY

Observability refers to measuring a system's internal states by examining its outputs. It encompasses three main pillars: logs, metrics, and traces. By leveraging logs, metrics, and traces, engineers gain deep insights into the internal workings of their applications, enabling them to diagnose issues, monitor health, and ensure reliable operation. Each of these components plays a crucial role in providing a comprehensive view of the system's behavior.

#### A. Logs

Logs are detailed, time-stamped records of events that occur within a system. They provide a granular, sequential history of what has happened, such as errors, warnings, and informational messages. Logs are crucial for diagnosing issues and understanding system behavior, as they offer context-rich information that can be analyzed post-incident to determine the root cause of a problem. The study found that organizations leveraging advanced observability practices, including log analytics, experienced a 69% reduction in mean time to resolution (MTTR) for unplanned downtime [3].



# International Journal for Research in Applied Science & Engineering Technology (IJRASET) ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538

Volume 12 Issue IX Sep 2024- Available at www.ijraset.com

#### B. Metrics

Metrics are numerical data points that provide insights into the performance and health of a system over time. They typically include statistics like CPU usage, memory consumption, request rates, and error rates. Metrics help in monitoring the system's overall state and are essential for identifying trends, detecting anomalies, and triggering alerts when certain thresholds are breached, enabling proactive system management. The IDC study revealed that organizations with mature observability practices saw a 66% faster MTTR for customer-impacting incidents [3].

#### C. Traces

Traces track the flow of requests as they traverse through various components of a distributed system, capturing the path and timing of each operation. They are vital for understanding the dependencies and performance bottlenecks within complex architectures, especially in microservices environments. Traces allow engineers to visualize the entire lifecycle of a request, pinpointing where delays or failures occur in the system. The study highlighted that organizations with advanced observability practices experienced a 63% reduction in the frequency of outages [3].

The real power of observability lies in the combination of logs, metrics, and traces. Each pillar provides a different perspective, offering a holistic view of system health and performance. To illustrate the synergy of these three pillars, consider an online banking application experiencing high latency. Metrics would show increased response times, logs could reveal repeated timeout errors, and traces would pinpoint the exact delay point. This combined information allows engineers to identify and resolve the root cause quickly.

Understanding and implementing the three pillars of observability are crucial for maintaining system reliability, diagnosing issues efficiently, and optimizing performance. By leveraging these components, software architects can gain deep insights into their systems, enabling proactive issue detection and resolution. The study found that organizations with mature observability practices experienced 74% faster mean time to detection (MTTD) for application performance degradation [3].

Faster execution of business transactions

Reduced time to run batch processes

Improved application performance (lower latency)



Fig. 1: Performance-Related Benefits of Observability (% improvement) [3]

#### III. DESIGNING AN OBSERVABILITY STRATEGY

An effective observability strategy begins with clearly understanding the system's architecture and business requirements. A crucial step in this process is defining Key Performance Indicators (KPIs).

#### A. Defining Key Performance Indicators (KPIs)

Defining KPIs is a critical process that involves identifying specific, quantifiable measures that accurately gauge the performance of various aspects of your system. When defining KPIs for observability, consider the following steps:

- 1) Align with Business Objectives: Ensure that your KPIs reflect your organization's overall goals and priorities.
- 2) Identify Critical System Components: Determine which parts of your system are most crucial to monitor.
- 3) Choose Relevant Metrics: Select metrics that provide meaningful insights into your system's health and performance.
- 4) Set Thresholds: Establish baseline performance levels and acceptable ranges for each KPI.
- 5) Ensure Measurability: Make sure your chosen KPIs can be accurately and consistently measured.
- 6) Consider Different Perspectives: Include KPIs that cover various aspects of your system, such as user experience, system performance, and business impact.

Examples of commonly used KPIs in observability strategies include:

- a) Latency: Measure the time it takes for requests to be processed.
- This can be broken down into various components:
- Network latency: Time taken for data to travel across the network.
- > Application latency: Processing time within the application.
- > Database latency: Time taken for database queries and responses.



ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538 Volume 12 Issue IX Sep 2024- Available at www.ijraset.com

- Best practices often suggest aiming for sub-second response times, with many e-commerce sites targeting under 100ms for critical operations.
- Tools like Prometheus or Datadog can be used to monitor and alert on latency issues.
- b) Error Rates: Track the frequency of errors within the system.
- This can include:
- ➢ HTTP error codes (e.g., 4xx client errors, 5xx server errors)
- > Application-specific errors (e.g., failed transactions, timeouts)
- > Infrastructure errors (e.g., out of memory errors, disk full errors)
- Many organizations aim to keep error rates below 0.1% for critical services.
- Error rates can be monitored using tools like ELK stack (Elasticsearch, Logstash, Kibana) or Splunk.
- c) Throughput: Monitor the number of transactions processed over time.
- This can be measured in various ways:
- Requests per second (RPS)
- Transactions per minute (TPM)
- Orders processed per hour
- Throughput KPIs should be set based on business requirements and system capabilities.
- Tools like Apache JMeter or Gatling can be used for load testing to determine system throughput capacity.
- *d) Resource Utilization: Observe CPU, memory, and storage usage.*
- This includes:
- > CPU usage: Percentage of CPU capacity being used.
- Memory usage: Amount of RAM being consumed.
- > Disk I/O: Rate of read/write operations on storage devices.
- > Network usage: Bandwidth consumption.
- Best practices often suggest keeping CPU utilization below 70-80% to allow headroom for traffic spikes.
- Cloud providers offer native monitoring tools (e.g., AWS CloudWatch, Azure Monitor) for resource utilization tracking.

When implementing these KPIs, it's important to:

- > Set appropriate thresholds based on your specific system and business needs.
- > Implement alerting mechanisms when KPIs exceed defined thresholds.
- Regularly review and adjust KPIs as your system evolves and business requirements change.
- ▶ Use visualization tools (e.g., Grafana, Kibana) to create dashboards for easy monitoring of these KPIs.

By carefully defining your KPIs, you create a solid foundation for your observability strategy, enabling more effective monitoring and quicker resolution of issues.

# B. Establishing Data Collection Points

Once KPIs have been defined, the next essential step is to establish data collection points that will capture the necessary metrics, ensuring that you gather the right information to monitor and analyze your system's performance effectively. Determine where to collect data within your system to ensure comprehensive coverage:

- 1) Application Code: Instrumenting application code is a fundamental aspect of observability. It involves strategically adding code to your application to emit valuable data:
- *a)* Logs: Insert logging statements at key points in your code to record events, errors, and important state changes. For example, log user actions, database queries, or API calls.
- *b)* Metrics: Implement counters, gauges, and histograms to measure specific aspects of your application's behavior. This could include request counts, response times, or business-specific metrics like order values.
- *c)* Traces: Add trace spans to track the flow of requests through your application. This is particularly important in microservices architectures to understand how requests propagate across different services.

Best practices include using structured logging formats (like JSON), consistent naming conventions, and appropriate log levels. Many languages have libraries (e.g., OpenTelemetry) that simplify this instrumentation process.



# International Journal for Research in Applied Science & Engineering Technology (IJRASET) ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538

Volume 12 Issue IX Sep 2024- Available at www.ijraset.com

- *Infrastructure:* Monitoring infrastructure involves collecting data from the underlying systems that support your application:
- Servers: Monitor CPU usage, memory consumption, disk I/O, and network traffic. This applies to both physical and virtual servers.
- > Containers: In containerized environments like Kubernetes, monitor container health, resource usage, and orchestration metrics.
- Network Components: Track network latency, throughput, and error rates. This includes routers, load balancers, and other network devices.
- > Databases: Monitor query performance, connection pools, and storage metrics.

Tools like Prometheus with node\_exporter, cAdvisor for containers, or cloud-native solutions like AWS CloudWatch or Azure Monitor can be used to collect these metrics.

- External Dependencies:
- Modern applications often rely on various third-party services and APIs. Monitoring these is crucial for understanding your system's overall health:
- API Monitoring: Track the performance and availability of external APIs. This includes response times, error rates, and data quality.
- Synthetic Monitoring: Use automated tests to simulate user interactions with your system and its dependencies. This can help detect issues before they impact real users.
- Service Level Agreements (SLAs): Monitor whether external services are meeting their promised SLAs.
- > Dependency Mapping: Maintain an up-to-date map of all external dependencies to understand their impact on your system.

External dependencies can be monitored using tools like Postman for API testing, Pingdom for synthetic monitoring, or more comprehensive solutions like Dynatrace.

Collecting data from these three key areas—application code, infrastructure, and external dependencies can help you build a comprehensive view of your entire system. This holistic approach to observability enables faster troubleshooting, proactive issue detection, and a deeper understanding of your application's behavior in production.

Ensure that data collection is efficient and does not introduce significant overhead. Use asynchronous logging and non-blocking instrumentation to minimize performance impact. Some advanced tracing systems add less than 1% overhead to request processing. When implementing these strategies, it's crucial to consider data privacy and security. The study highlights that organizations with mature observability practices are 2.9 times more likely to identify issues before they impact customers, emphasizing the importance of a well-designed, secure observability strategy [3].



Fig. 2: Business KPIs (% improvement) [3]

## C. Choosing the Right Tools

Once the data collection points are clearly established, the next critical step is to choose the right observability tools that can effectively gather, analyze, and visualize the data, ensuring that you can monitor your system's performance in alignment with your defined KPIs. Consider the following categories:

## 1) Logging Tools:

- ELK Stack (Elasticsearch, Logstash, Kibana): This popular open-source stack offers a comprehensive logging solution. Logstash collects and processes logs from multiple sources, Elasticsearch indexes and stores the data, and Kibana provides a user-friendly interface for visualization and analysis.
- Fluentd: An open-source data collector that unifies data collection and consumption, making it easy to collect, process, and forward logs to various backends, including Elasticsearch.



These logging tools are crucial for centralized log management, allowing teams to aggregate logs from diverse sources, search through them efficiently, and visualize log data for easier analysis and troubleshooting.

## 2) Metrics Tools:

- Prometheus: An open-source monitoring and alerting toolkit designed for reliability and scalability. It's particularly well-suited for dynamic environments like Kubernetes and features a powerful query language (PromQL) for data analysis.
- Graphite: A scalable, real-time graphing system that lets you collect, store, and visualize time-series data. It's known for its simplicity and ability to handle large amounts of data.

These tools are essential for collecting and analyzing numerical data about your system's performance, helping teams identify trends, anomalies, and potential issues before they become critical.

## 3) Tracing Tools:

- Jaeger: An open-source, end-to-end distributed tracing system that helps monitor and troubleshoot transactions in complex distributed systems. It's particularly useful in microservices architectures.
- Zipkin: Another open-source distributed tracing system that helps gather timing data needed to troubleshoot latency problems in service architectures. It manages both the collection and lookup of this data.

Tracing tools are vital for understanding the flow of requests through complex, distributed systems, helping teams identify performance bottlenecks and optimize system behavior.

## 4) Visualization Tools:

- Grafana: An open-source platform for monitoring and observability. It allows you to query, visualize, alert on, and understand your metrics no matter where they are stored. It's particularly powerful for creating comprehensive dashboards.
- Kibana: Part of the ELK stack, Kibana lets you visualize your Elasticsearch data and navigate the Elastic Stack. It's great for log and time-series analytics, application monitoring, and operational intelligence use cases.

These visualization tools are crucial for creating intuitive, informative dashboards that help teams quickly grasp the state of their systems and identify issues at a glance.

Each tool plays a specific role in the observability ecosystem, and many organizations use a combination of them to build a comprehensive observability strategy. The choice of tools often depends on factors like the organization's specific needs, the existing technology stack, scalability requirements, and team expertise.

Emerging tools and platforms increasingly leverage AI and machine learning to provide predictive insights and automated anomaly detection. These advanced tools can help identify patterns and predict potential issues before they occur. According to Gartner's Market Guide for AIOps Platforms, by 2024, 30% of large enterprises will use AI-augmented observability tools to reduce MTTR by 50% [4].



Key Considerations for Choosing Observability and AIOps Tools [5]



ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538 Volume 12 Issue IX Sep 2024- Available at www.ijraset.com

## IV. IMPLEMENTING OBSERVABILITY

Implementing observability involves putting into practice the strategies and tools that enable comprehensive monitoring and analysis of your system's performance. This step is where the theoretical planning turns into actionable measures, ensuring that your systems are not only monitored but also provide meaningful insights that drive reliability and efficiency.

#### A. Logging

Effective logging involves:

- 1) Structured Logging:
- Practice: Use a consistent format (e.g., JSON) to make logs machine-readable. This aids in automated parsing and analysis, allowing for easier search and correlation of log entries.
- Implementation: Embed logging statements in your code at critical points, such as error handling, key operations, and external service calls. Use structured logging to include contextual information like request IDs, user IDs, and timestamps.
- Impact: According to the Splunk State of Observability Report, organizations using structured logging have significantly improved their ability to troubleshoot issues quickly [5].
- 2) Log Levels:
- Practice: Differentiate logs by severity (e.g., DEBUG, INFO, WARN, ERROR). This helps filter and focus on relevant issues during debugging and incident investigation.
- Implementation: Define and consistently apply log levels across your application. Configure logging frameworks to emit appropriate levels based on the environment (e.g., more verbose in development, less in production).
- Recommendation: The report recommends using a limited number of log levels to maintain clarity and ease of use [5].
- 3) Centralized Logging:
- Practice: Aggregate logs from various services into a centralized system to facilitate easy searching and correlation of events.
- Implementation:
- Select and configure log aggregation tools (e.g., Fluentd, Logstash) to collect logs from different services.
- > Set up parsing rules to extract relevant information from log entries.
- > Configure the tools to forward logs to a centralized storage system (e.g., Elasticsearch).
- > Utilize visualization tools (e.g., Kibana) to create dashboards, set up saved searches, and monitor key logs.
- Impact: The Splunk report indicates that organizations using centralized logging have reported reduced Mean Time to Resolution (MTTR) [5].
- 4) Continuous Improvement:
- Practice: Regularly review and refine your logging practices.
- Implementation:
- Conduct periodic reviews of logging effectiveness and adjust as needed.
- > Optimize log retention policies and storage strategies.
- > Provide ongoing training to team members on effective log analysis and troubleshooting techniques.

#### B. Metrics

Metrics are crucial for monitoring system performance and health. Here's a comprehensive guide to implementing metrics in your observability strategy:

- 1) Metric Definition and Instrumentation:
- Define metrics that align with your Key Performance Indicators (KPIs), such as request counts, latency, and error rates.
- Integrate metrics collection within your codebase using libraries (e.g., Micrometer for Java).
- Add code to emit metrics at critical points, such as request counts, response times, and error rates.

2) Metric Collection and Storage:

- Configure tools (like Prometheus) to scrape metrics endpoints at regular intervals.
- Set up efficient data storage systems to handle the volume of metrics data.
- Implement data retention policies to balance storage costs with data availability.
- 3) Metric Analysis:
- Use powerful query languages (like PromQL for Prometheus) to extract insights from the collected data.



ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538 Volume 12 Issue IX Sep 2024- Available at www.ijraset.com

- Set up alerting rules based on metric thresholds to identify issues proactively.
- 4) Visualization:
- Use visualization tools like Grafana to create informative dashboards.
- Design dashboards that provide real-time insights and historical trends.
- Ensure dashboards allow for filtering and drilling down into specific periods or components.
- Create panels for key metrics and set up alerts for significant changes or anomalies.
- 5) Continuous Improvement:
- Regularly review the effectiveness of your metrics and adjust as needed.
- Stay updated with new metrics libraries and tools to enhance your observability capabilities.

## C. Tracing

Tracing involves:

- 1) Trace Instrumentation:
- Instrument your code to generate trace spans using tracing libraries.
- Add trace spans at critical points, such as incoming requests, external service calls, and database queries.
- 2) Trace Context Propagation:
- Ensure trace identifiers (trace IDs and span IDs) are propagated through all services.
- Use tracing libraries that support automatic context propagation or manually pass trace context in headers.
- Adhere to industry standards for propagating context across service boundaries.
- 3) Trace Collection and Storage:
- Set up a backend system to collect and store trace data.
- Configure your services to send trace data to the tracing backend.
- 4) Trace Analysis and Visualization:

Analyze trace data to identify performance issues and dependencies between services.

- Use visualization tools to identify latency hotspots and bottlenecks.
- Create dashboards that help in understanding the flow of requests and pinpointing where time is spent.

# V. BEST PRACTICES FOR OBSERVABILITY

- Start with Business Goals: Align observability metrics with business objectives to ensure they provide actionable insights. According to the Dynatrace Global CIO Report, 90% of CIOs say their organizations need to improve the link between observability and business outcomes to maximize cloud investments [3]. Monitor metrics related to user experience, such as response times and error rates, especially if customer satisfaction is a key goal.
- 2) Automate Where Possible: Use automation to ensure consistent and comprehensive monitoring, reducing the risk of human error. Automate the deployment and configuration of observability tools and use automated alerts. The Dynatrace report reveals that 92% of CIOs say AI-powered observability capabilities are key to accelerating innovation while minimizing risk [3].
- *3)* AI and Predictive Analytics: Incorporate AI-driven observability tools to enhance monitoring capabilities. 95% of CIOs say AI-powered observability is essential for reducing digital transformation risks and accelerating innovation [3].
- 4) Alerting and Incident Response: Set up alerting mechanisms to detect issues promptly and define incident response procedures to address them effectively. Use tools to configure alerts based on key metrics and set up escalation policies to ensure timely response.
- 5) Security and Compliance: Incorporate security best practices into your observability strategy. Ensure sensitive data is properly masked or encrypted in logs and traces. The IBM Cost of a Data Breach Report 2023 found that the average data breach cost reached \$4.45 million in 2023, emphasizing the need for robust security measures [6].
- 6) End-to-end Observability: Strive for comprehensive observability across your entire technology stack, including applications, infrastructure, and external dependencies.
- 7) Contextual Correlation: Ensure that logs, metrics, and traces can be correlated across your system. This allows for faster rootcause analysis during incidents.



## International Journal for Research in Applied Science & Engineering Technology (IJRASET) ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538

Volume 12 Issue IX Sep 2024- Available at www.ijraset.com

- 8) Data Retention and Scalability: Plan for data retention and scalability from the outset. As systems grow, the volume of observability data can increase exponentially. Implement data lifecycle management policies to balance data retention needs with storage costs.
- 9) Training and Culture: Foster a culture of observability within your organization. Train developers and operations teams to use observability tools and interpret the data effectively.
- 10) Regularly Review and Iterate: Improve your observability strategy based on feedback, new requirements, and technological advancements. Review your observability setup regularly and make adjustments to ensure it remains effective.
- 11) Synthetic Monitoring: Implement synthetic monitoring to detect issues before they impact real users proactively.

## VI. CONCLUSION

As software systems grow in complexity, observability has become a critical component of modern architecture. This article has outlined key strategies and best practices for designing and implementing effective observability solutions, emphasizing the importance of a holistic approach that combines logs, metrics, and traces. By aligning observability with business objectives, leveraging automation, and adopting emerging technologies such as AI and machine learning, organizations can significantly improve their ability to monitor, troubleshoot, and optimize their systems. The future of observability lies in more proactive, predictive approaches that not only react to issues but anticipate and prevent them. As the field evolves, continuous learning and adaptation will be crucial for software architects and engineers to stay ahead of the curve and ensure the reliability and performance of their systems in an increasingly digital world.

#### REFERENCES

- [1] Gartner, "Market Guide for Digital Business Observability," Gartner, Inc., Tech. Rep., 2023. [Online]. Available: <u>https://www.gartner.com/</u> <u>en/documents/5533895</u>
- [2] New Relic, "2023 Observability Forecast," New Relic, Inc., Tech. Rep., 2023. [Online]. Available: <u>https://newrelic.com/observability-forecast/2023/state-of-observability</u>
- [3] IDC, "The Business Value of Observability," sponsored by New Relic, April 2022. [Online]. Available: <u>https://newrelic.com/sites/default/files/2022-04/new-relic-idc-bv-white-paper-%23us48924422-2022-04-27.pdf</u>
- [4] Gartner, "Market Guide for AIOps Platforms," Gartner, Inc., Tech. Rep., 2023. [Online]. Available: <u>https://www.gartner.com/en/documents/4000217</u>
- [5] Splunk, "The State of Observability 2023," Splunk Inc., Tech. Rep., 2023. [Online]. Available: <u>https://www.splunk.com/en\_us/form/state-of-observability.html</u>
- [6] IBM, "Cost of a Data Breach Report 2023," IBM Security, Tech. Rep., 2023. [Online]. Available: https://www.ibm.com/reports/data-breach











45.98



IMPACT FACTOR: 7.129







INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 🕓 (24\*7 Support on Whatsapp)