


WHITE PAPER



The Proactive Edge

Mastering Observability for Operational Resilience

*A Comprehensive Guide to
Full Stack Observability*



Embracing Observability for Enhanced IT Performance

Executive Summary

This white paper serves as an introduction to the significance and benefits of Observability in dynamic IT environments. By exploring Observability's role in monitoring, analysing, and optimising system performance, this document aims to explain how Observability empowers organisations to achieve operational excellence.

Throughout this white paper, we will deep dive into the core principles of Observability, like real-time monitoring, types of Observabilities, Application Performance Monitoring Tools (APM Tools), and proactive troubleshooting. Additionally, we will explore the concepts of Data Observability, Applied Observability, and Full Stack Observability (FSO) and its transformative potential in optimising resource utilisation and ensuring seamless system functionality.



Join us on this journey as we uncover the advantages of Observability and FSO, offering insights into their application in the modern IT landscapes and their impact on organisational success.



Table of Contents

● <u>What is Observability?</u>	4
● <u>Why is Observability Important?</u>	5
● <u>Difference between Monitoring and Observability</u>	6
● <u>Types of Observability</u>	7
● <u>The Observability Edge</u>	8
● <u>Security Boost: Observability's Real-Time Vigilance in SD-WAN and SASE</u>	10
● <u>Evolution of Observability: From Observability to Full Stack Observability</u>	12
● <u>Full Stack Observability: Key Components</u>	13
● <u>Benefits of Full Stack Observability</u>	16
● <u>Implementing Full Stack Observability: Best Practices</u>	17
● <u>Pioneering Full Stack Observability: Conclusion</u>	18
● <u>About Cisco AppDynamics and Cisco ThousandEyes</u>	19
● <u>About Proactive Data Systems Private Limited</u>	19

What is Observability?

An Overview



Just as a symphony's conductor blends a host of musical instruments into a brilliant masterpiece, a multitude of elements work in sync in the modern, interconnected digital ecosystem to deliver cutting-edge solutions. The fusion of Cloud, Software-Defined Wide Area Networking (SD-WAN), Secure Access Service Edge (SASE), Applications, and Networks has ushered in an era of digital complexity. The complexity is further amplified in native cloud environments, where infrastructure and distributed services require advanced monitoring and management strategies. This complexity demands organisations to scale up their digital game for better operational efficiency, and enhanced customer experience.

Within this intricate web of technologies understanding what is happening inside the systems has become critical. It's really important to know how these complex systems work, and that's what we mean when we talk about Observability. It's a clear picture of what's going on inside the many layers of our technology stack so that we can keep everything running with surgical precision.



Let us understand Observability through a simple example. Imagine you're driving a car. When everything is functioning smoothly, you're able to see your speedometer, fuel gauge, battery level indicator, and engine temperature, and you can hear the engine running smoothly. If there's an unusual sound or a warning light, you know something is wrong and demands immediate attention. This ability to observe your car's internal state through a set of external alerts and alarms is similar to what Observability offers in the world of IT.

Observability helps us understand how a system works by observing it from the outside. Just like checking a car's gauges

ensures a smooth ride, observability tools use data from sources like metrics, events, logs, and traces (MELT data) to monitor software, networks, servers, and other parts of the system, ensuring everything runs like clockwork and problems are quickly spotted.

The significance of observability in complex IT environments is underscored by a report published by Gartner in October 2022. In this report, Applied Observability ranks at number two in the list of *Top 10 Strategic Technology Trends for 2023*,¹ emphasising its critical role in shaping the future of IT operations and management.

Why is Observability Important?



Observability plays an important role in ensuring the reliability, performance, and resilience of IT systems. It enables organisations to proactively identify and address potential problems before they escalate, optimise system performance, and deliver exceptional user experiences.

A Gartner article *Monetising Observable Data Will Separate the Winners and Losers*² published in November 2022 predicts that by 2026, 70% of organisations leveraging observability will achieve faster decision-making, giving them a competitive edge in key business or IT areas. The study further reveals that using observability can speed up response times and improve business operations instantly.

A survey conducted in India by the premier market intelligence firm IDC underscores this fact. The IDC *Future of Digital Infrastructure 2022 Global Sentiment Survey*³ in India revealed that 42.5% of participants believe that Observability, combined with Artificial Intelligence/Machine Learning (AI/ML) will significantly improve infrastructure operations and enhance application performance.

Difference between Monitoring and Observability

Monitoring is Reactive

Observability is Proactive

Monitoring helps you fix a problem after it has occurred. Monitoring systems are designed to detect and alert you to problems or issues in your system or processes after they have already occurred or crossed predefined thresholds. Monitoring is reactive because it involves setting thresholds or alerts based on predefined conditions. When those conditions are met or exceeded, the monitoring system reacts by triggering alerts or notifications. It focuses on responding to known patterns or deviations from expected norms.

Observability, on the other hand, enables you to monitor system behaviour in real-time, allowing you to identify early signs of potential issues before they escalate. Unlike monitoring, which is reactive and responds to predefined conditions, observability provides a comprehensive view of system performance in real time.

Observability enables proactive problem-solving by identifying potential issues before they escalate. It enables understanding complex interactions, and optimising system performance based on continuous analysis and understanding. Observability involves collecting and analysing diverse data sources, such as logs, metrics, traces, and events, to gain insights into how different components of the system work together.

Observability enables IT teams to identify complex dependencies and uncover hidden issues by deep diving into the root causes of problems. By focusing on proactive problem-solving, observability empowers organisations to improve overall system resilience.

While both Observability and Monitoring are essential for maintaining system health, they serve distinct purposes and offer unique insights. Let's illustrate the difference between monitoring and observability through a simple example of a delivery service.

Monitoring Use-Case Example:

Imagine you run a delivery service where your primary goal is to ensure packages are delivered on time. To monitor delivery operations, you track metrics like delivery times, packages delivered per hour, and availability of drivers. If a delivery takes longer than expected, your monitoring system sends an alert, allowing you to take corrective action such as reassigning drivers or sending regret emails to customers informing them about delays. That is how Monitoring works - it focuses on tracking specific metrics and reacting to predefined thresholds or patterns.

Observability Use-Case Example:

Now, let's consider observability in the context of the same delivery service. Observability goes beyond just tracking metrics. It involves taking a holistic view of the entire delivery process, including

the interactions between different components like drivers, vehicles, routes, and customer locations. With observability, you collect data not only on delivery times but also on factors like traffic conditions during specific periods, driver behaviour, fuel economy, package handling, and customer feedback.

For example, suppose you notice a sudden increase in delivery delays. With observability, you can dive deep into the data to identify the root causes. You might discover that certain routes have high traffic congestion during specific times of the day, leading to delays. You might also discover that some drivers are consistently taking longer breaks, impacting delivery timelines. Armed with these insights, you can take proactive steps like adjusting routes, optimising schedules, reviewing delivery deadlines, or providing additional training to drivers.

Types of Observability

- **Observability**

As stated above, this covers monitoring, understanding, and managing complex systems by collecting and analysing data from sources like logs, metrics, traces, and events.

- **Data Observability**

Specifically focuses on monitoring and ensuring data quality, reliability, and accessibility within the system. Data Observability solutions help organisations monitor data flows, detect anomalies, analyse data quality metrics, and ensure data governance compliance.

- **Applied Observability**

Involves implementing observability principles and tools to improve system performance and reliability within an organization's IT infrastructure. Applied Observability solutions, on the

other hand, focus on monitoring application performance metrics, analysing system logs, providing real-time insights into network performance, and integrating with automation tools for proactive troubleshooting.

The Observability Edge

The substantial advantages that Observability offers over traditional Monitoring are driving a growing number of companies to adopt advanced Observability solutions. A research report titled *The Age of Application Observability* by AppDynamics highlights this trend, with 78% of technologists noting that the escalating data volumes from multi-cloud and hybrid environments are rendering manual monitoring impractical.

The findings were derived from interviews conducted with 1,140 IT professionals working in organisations with a minimum turnover of \$500 million. These professionals were located in 13 countries, including India, Australia, Brazil, Canada, Colombia, France, Germany, Japan, Mexico, Singapore, the United Arab Emirates, the United Kingdom, and the United States.

The report also reveals that 85% of technologists consider observability as a strategic priority for their organisations. This emphasises the growing recognition of observability's importance in modern IT operations and decision-making processes.




The Age of Application Observability by AppDynamics report reveals that **85%** of technologists consider observability as a strategic priority for their organisations

The report highlights that transitioning to cloud technologies is increasing the intricacy of the IT landscape, necessitating the use of advanced Observability tools to manage the complexities, particularly in native cloud environments. According to the report, 83% of technologists agree that the shift to cloud-native technologies is causing a rise in complexity, mainly due to microservices and containers creating a large amount of MELT (metrics, events, logs, and traces) data.

Monitoring is reactive, responding to known patterns or threshold
Observability is proactive, enabling real-time analysis to identify emerging anomalies

**Here's a table listing the key differences between
Monitoring and Observability:**

Aspect	Monitoring	Observability
Focus	Tracks specific metrics and indicators	Provides 360-degree understanding of system behaviour
Data Sources	Metrics, thresholds, alerts, dashboards	Logs, metrics, traces, events
Purpose	Ensures system health and performance	Enables deep analysis, troubleshooting, and proactive problem-solving
Reactivity	Reactive - triggers actions on breach of predefined thresholds or patterns	Proactive - goes beyond reactive alerts. Allows exploration and understanding even in unknown scenarios
Visibility	Limited visibility into internal system workings	Comprehensive visibility, even into complex interactions and dependencies
Tools and Techniques	Monitoring tools, thresholds, alerting mechanisms	Observability platforms, distributed tracing, log aggregation
Root Cause Analysis	Limited root cause analysis based on monitored metrics	Advanced root cause identification through comprehensive data analysis
Scalability	Good for tracking known patterns and metrics	Better suited for scaling and handling complex and dynamic IT environments
Use Cases	Performance monitoring, resource utilisation tracking	Troubleshooting complex issues, optimising system behavior



The Name is O11Y!

FUN FACT

Developers often refer to Observability as O11Y, pronounced as OLLY, to simplify its pronunciation. This abbreviation was derived by taking the first letter "O" from "observability" and the last letter "Y," with 11 letters between them, resulting in the term O11Y. This creative naming convention has become a common shorthand in the tech industry for expressing the concept of observability.

Security Boost: Observability's Real-Time Vigilance in SD-WAN and SASE

As more and more businesses scramble to embrace digital transformation and adopt modern networking architectures such as SD-WAN and SASE, robust observability becomes paramount. Just like any other complex IT ecosystem, Observability plays an important role in providing comprehensive visibility into the performance, security, and health of SD-WAN and SASE ecosystems.



Comprehensive Visibility

SD-WAN and SASE architectures introduce complexities with distributed networks, cloud connectivity, and diverse endpoints. Observability solutions offer comprehensive visibility into these dynamic and interconnected components, allowing IT teams to monitor network traffic, application performance, user behaviour, and security events across the entire infrastructure.



Proactive Monitoring

With observability, organisations can proactively monitor key metrics and KPIs related to SD-WAN and SASE performance. This includes monitoring network latency, bandwidth utilisation, application response times, user access patterns, security incidents, and compliance adherence.



Deep Insights and Analysis

Observability correlates data from various sources within SD-WAN and SASE environments including correlating network telemetry data, security logs, application performance metrics, and user activity logs to gain a complete understanding of system behaviour. These insights enable IT teams to perform root cause analysis and troubleshoot complex issues.



Enhancing Security Posture

Observability offers real-time visibility into network and application performance, allowing IT teams to detect anomalies, bottlenecks, or security threats. Real-time alerts and notifications enable swift response to incidents, reducing Mean Time to Detect (MTTD) and Mean Time to Resolve (MTTR), thereby improving overall security posture.




Optimising Resources

Observability helps in optimising resource utilisation within SD-WAN and SASE environments by monitoring resource consumption, traffic patterns, and workload distribution. Organisations can identify underutilised resources, optimise network paths, and dynamically allocate resources based on demand.

Evolution of Observability: From Observability to Full Stack Observability (FSO)

“ *It is not the strongest of the species that survive, nor the most intelligent, but the one most responsive to change*

- Charles Darwin



Limitations of Observability

Alas, even the most brilliant innovations bear the weight of limitations, and Observability is no exception. The scope of Observability often remains confined to individual components or layers within the entire IT stack. This narrow focus can pose challenges in achieving a birds-eye view of how these components interact and influence the overall system performance. Without a 360-degree view of system dynamics and interdependencies, IT teams may find it difficult to identify the root causes of complex issues or optimise system-wide performance effectively.

For example, while Observability may reveal issues with a specific application's response time or a server's CPU utilisation, it may not provide insights into how these issues are interconnected or how they affect end-to-end system functionality.

This limited visibility across the entire IT stack hinders the ability to optimise system-wide performance, and proactively address potential issues before they escalate.



Enter Full Stack Observability

To bridge this gap, organisations are increasingly turning to Full Stack Observability (FSO). FSO extends the principles of Observability across all layers and components of the IT environment, including applications, infrastructure, networks, and services. This approach enables IT teams to gain a unified view of

system behaviour, identify complex dependencies and interactions, and proactively optimise performance and reliability across the entire IT stack. By transitioning from traditional Observability to full stack Observability, organisations can unlock deeper insights, and enhance efficiency.

Full Stack Observability: Key Components

FSO encompasses several essential elements and technologies that work in sync to provide a comprehensive view of system behaviour across all components of the IT ecosystem. This includes applications and the infrastructure. Let's break down these key components:

- **Logs**

Logs are chronological records of events or activities generated by various components within the IT infrastructure, such as apps, servers, databases, and networking devices. Logs capture valuable information about system operations, errors, user activities, and security events. Analysing logs helps in troubleshooting issues, detecting anomalies, and monitoring system health. Think of logs as the detectives of IT, revealing the 'What' of IT processes with every line of code and timestamp behind every system event and action, piece by piece.

Logs are like detailed diaries kept by systems, meticulously recording every event, error, and action that occurs within the system. They serve as a historical record, providing a chronological account of system activities, errors, warnings, and informational messages.

Logs capture the 'What' of system behaviour, offering insights into events such as user interactions, application processes, security incidents, and system events. Analysing logs helps in understanding the sequence of events, identifying patterns, troubleshooting issues, and auditing system activities for compliance and security purposes.

- **Metrics**

Metrics are quantitative measurements that provide insights into the performance and health of IT systems. These measurements include CPU utilisation, memory usage, disk Input/Output (I/O), network traffic, response times, and other key performance indicators (KPIs). Metrics help in monitoring system performance, identifying trends, setting thresholds for alerts, and optimizing resource utilisation.

Metrics help us delve deeply into the 'What' highlighted by logs. They act as mathematical storytellers, conveying the 'What' of system health, efficiency, and effectiveness through numerical insights. By monitoring and analysing metrics, IT teams can gain visibility into system health, identify trends, proactively address performance issues, and make data-driven decisions to improve overall system efficiency and effectiveness.

- **Traces**

Traces are records of transactions or activities that span multiple components or services within a distributed system. Tracing allows IT teams to trace the path of a request or transaction across various microservices, containers, APIs, and network layers. Traces provide visibility into dependencies, latency issues, and performance bottlenecks across distributed architectures.

Traces are the digital footprints that unveil the 'How' of system interactions, allowing us to zoom in on the details and connections within the system. They provide a granular view of how individual components interact and communicate with each other during system operations.

By following the path of data and requests across various services, microservices, APIs, and network layers, traces unveil the step-by-step journey of a transaction or process.

This detailed insight into the 'How' of system behaviour enables IT teams to understand the root causes of issues, identify performance bottlenecks, and troubleshoot complex problems effectively. Traces help in pinpointing where and why an issue occurred, facilitating quick resolution and optimisation of system performance.

- **Events**

Events are key signals that tell us what's happening in our systems in real-time. They're like alerts that notify us about various activities and incidents, such as application crashes, security breaches, changes in settings, user actions, system updates, and more.

Keeping a close eye on events is important for staying ahead of problems and responding quickly. Just like we monitor logs, metrics, and traces, monitoring events is a vital part of FSO strategies. It helps us spot issues early and manage incidents effectively.

- **Distributed Tracing**

Distributed tracing is like a digital detective, following the path of transactions as they journey through interconnected services and components within a distributed system. It employs advanced techniques to capture and correlate trace data from various services,

reconstructing the complete end-to-end transaction flow. By leveraging distributed tracing tools, IT teams can gain a visual representation of transaction paths, pinpoint performance bottlenecks, and fine-tune system performance in complex distributed environments.

- **Log Aggregation**

Log aggregation is like the conductor of a symphony, orchestrating the collection, storage, and analysis of logs from diverse sources into a centralised repository. Log aggregation tools harmonise logs from applications, servers, databases, network devices,

and cloud platforms, presenting a unified view of log data. This centralised approach to log aggregation streamlines log management, improves searchability, fosters event correlation, and supports proactive monitoring and troubleshooting efforts.

By leveraging these key components, organisations can achieve effective Full Stack Observability implementation, gain deep insights into system behaviour, detect, and resolve issues proactively.

Benefits of Full Stack Observability

Realising Business Impact: Full Stack Observability in Action

FSO brings a host of tangible benefits to IT operations, revolutionising the way organisations monitor their systems. Some of the key benefits of FSO include:

1

Improved Troubleshooting

FSO provides comprehensive visibility into the entire IT stack, allowing IT teams to pinpoint issues quickly and accurately. By correlating the MELT data across all layers, FSO enables efficient root cause analysis and troubleshooting, thus reducing downtime.

2

Faster Incident Resolution

With FSO, IT teams can proactively detect anomalies, and predict potential issues before they impact users. This approach speeds up the incident resolution and minimises service disruptions.

3

Optimised Resource Utilisation

By analysing metrics and performance data from the entire IT stack, FSO helps IT teams identify underutilised resources, fine-tune configurations, and make data-driven decisions to improve efficiency and reduce costs.

4

Enhanced System Performance

FSO empowers organisations to monitor KPIs in real time and identify opportunities for performance optimisation. By continuously analysing system behaviour and trends, FSO enables performance tuning, leading to enhanced responsiveness and reliability.

As organisations continue to embrace FSO and harness its potential, they are poised to realise significant business impact, ensuring resilient and high-performing IT environments that support their strategic objectives and drive success in the dynamic digital landscape.

Implementing Full Stack Observability

Best Practices

Selecting the Right Tools and Platforms

Implementing FSO requires careful consideration of various factors to ensure successful adoption. Here are some best practices for organisations looking to implement FSO:

Selecting the Right FSO Tools and Platforms

- Evaluate FSO tools and platforms based on their capabilities to collect, analyse, and visualise data from logs, metrics, traces, and events across the entire IT stack.
- Consider advanced solutions like AppDynamics and Thousand Eyes, renowned for their comprehensive observability features and ability to provide deep insights.

Monitoring Apps and Infrastructure

- Use monitoring agents, sensors, or APIs to collect data points for observability from applications and infrastructure components.
- Leverage instrumentation techniques offered by FSO tools to track performance metrics, capture logs, and generate traces for end-to-end visibility.

Integrating Data Sources

- Integrate data sources such as application logs, system metrics, network traffic data, and user interactions into a centralised observability platform.
- Ensure seamless data integration between different FSO tools and platforms to enable holistic monitoring and analysis across the IT stack.

Establishing Monitoring and Alerting Mechanisms

- Define KPIs and threshold values for monitoring critical metrics related to system health, performance, and availability.
- Set up proactive alerting mechanisms using FSO tools to receive real-time notifications about anomalies, deviations, or potential issues that require immediate attention.

Continuous Improvement and Optimisation:

- Implement a feedback loop for continuous improvement by analysing observability data, identifying areas for optimisation, and making data-driven decisions to enhance system performance.
- Regularly review and update observability strategies, tool configurations, and monitoring practices to adapt to evolving IT environments and business requirements.

By following these best practices and leveraging advanced FSO tools such as AppDynamics and Thousand Eyes, organisations can effectively implement Full Stack Observability and gain actionable insights to drive operational excellence, improve user experiences, and achieve business success.

Pioneering Full Stack Observability: Conclusion

In conclusion, Full Stack Observability has proven to be a game-changer for IT operations, offering tangible benefits that significantly impact outcomes. From improved troubleshooting capabilities and faster incident resolution to optimise resource utilisation and enhanced system performance, FSO revolutionizes the way organisations monitor and manage their IT ecosystems.

By providing comprehensive visibility into the entire IT stack and enabling proactive monitoring and alerting, FSO empowers IT teams to reduce downtime and enhance system reliability, make informed decisions, optimise resource allocation, and ultimately drive cost savings and operational efficiency.

About Cisco AppDynamics and Cisco ThousandEyes

Cisco AppDynamics and Cisco ThousandEyes excel in observability, specialising in application and network performance monitoring. AppDynamics grants insights into application behaviour, while ThousandEyes provides visibility into network paths and cloud services. Together, these solutions enable businesses to achieve comprehensive Full Stack Observability, proactively resolve issues, and deliver exceptional digital experiences.

Proactive Data Systems is an enterprise IT infrastructure and services company. For more than three decades, the company has designed, deployed, and operated the technology foundations behind some of India's most demanding enterprise environments, from regulated industries to large distributed operations.

Founded in 1991, Proactive serves more than 1,500 customers, from fast-growing businesses to Fortune 500 enterprises, across banking and financial services, manufacturing, healthcare, public sector, IT and ITES, retail, education, and global capability centres. Our work spans enterprise networking, cybersecurity, cloud and AI-ready data centres, collaboration platforms, structured cabling, and managed services.

A Cisco Preferred Partner under the Cisco 360 Partner Program, Proactive is recognised across the Networking, Security, Collaboration, Cloud & AI, and Services portfolios, and is a Cisco Advanced Customer Experience Specialised Partner. The company is ISO 9001:2015 and ISO/IEC 27001 certified, and engineers solutions in partnership with Dell, Nutanix, NetApp, Palo Alto Networks, Fortinet, VMware, Veeam, and CommScope.

Headquartered in Delhi NCR, with offices in Mumbai, Bengaluru, Pune, Hyderabad, Indore, and Singapore.

To learn more about our solutions, write to us
at pro@proactive.co.in

 www.proactive.co.in

PROACTIVE[™]
Enhancing Digital Experience