THE 5 STAGES OF THE OBSERVABILITY MATURITY MODEL

WHITE PAPER

Measure and enhance your business's ability to achieve active observability and beyond with the 5 stages of the observability maturity model by Apica.

www.apica.io

OVERVIEW

The 5 stages of the Observability Maturity Model by Apica include Level 0- **Monitoring**, Level 1- **Observability**, Level 2- **Active Observabilit**y, Level 3- **Intelligent Observability**, and Level 4- **Federated Observability**. Each level provides a more comprehensive understanding of the performance and behavior of systems and applications and allows organizations to quickly detect and resolve issues.

Table of Contents

Introduction

Monitoring has been around for quite a while now. It has been the conventional method to gain insights and monitor system performance. However, with the advent of observability, things got way more efficient.

Not that observability replaced monitoring but it is more of an augmentation/superset of it, if not a replacement.

Now, with the unprecedented expansion of data volumes in recent years, it has become difficult for businesses to keep up. Distributed systems have become particularly complicated due to their scale and complexity.

It has gotten extremely hard for DevOps, IT departments, and SREs to gather, combine, and analyze performance information on a large scale. Teams employ a wide array of techniques to discover the source of an issue, such as combining methods and tools or manually piecing together siloed data fragments. But traditional monitoring is time-consuming and does not provide acuity for how to improve business results.

Almost 65% of enterprises have more than 10 monitoring tools, many of which are used as silos to meet distinct requirements for different departments. Thus, observability is the next stage in monitoring evolution.

Observability Maturity Model

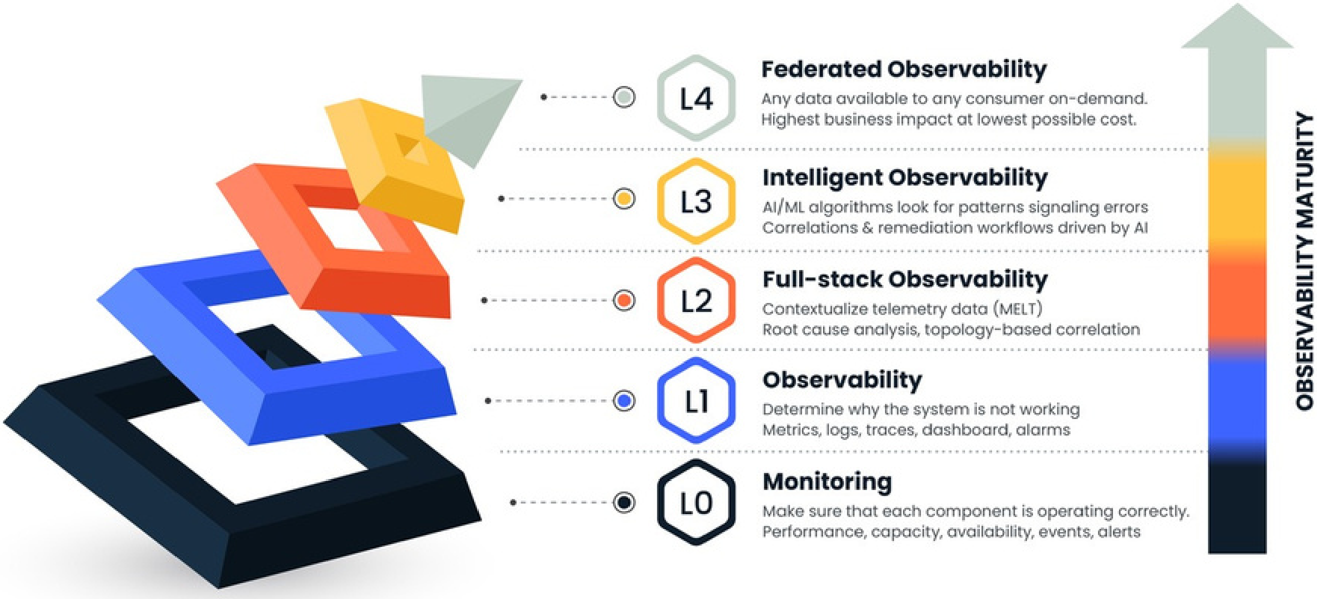

Federated Observability
Any data available to any consumer on-demand.
Highest business impact at lowest possible cost.
L4
Intelligent Observability
AI/ML algorithms look for patterns signaling errors
Correlations & remediation workflows driven by AI
L3
Full-stack Observability
Contextualize telemetry data (MELT)
Root cause analysis, topology-based correlation
L2
Observability
Determine why the system is not working
Metrics, logs, traces, dashboard, alarms
L1
Monitoring
Make sure that each component is operating correctly.
Performance, capacity, availability, events, alerts
L0
OBSERVABILITY MATURITY

In a Glimpse

Level 0- Monitoring:
Monitoring provides you with basic information regarding the health and status of individual components in your infrastructure, as well as warning you if one breaks down. Monitoring essentially gives insights into performance, capacity, availability, events, and alerts.

Level 1- Observability:
Observability is the degree to which an internal system's states may be deduced from external sources. Metrics, logs, and traces have traditionally been used as the three pillars of observable data.

Level 2- Active Observability:
Active observability goes beyond traditional monitoring by actively engaging with the systems and applications, providing a deeper understanding of their health and performance, and ultimately ensuring a smoother and more responsive digital experience for users.

Level 3- Intelligent Observability:
With intelligent observability or AIOps, advanced data analysis and machine learning techniques provide rapid insights into the performance and behavior of systems and applications. End users detect and resolve issues significantly faster compared to the other stages as the system learns patterns and correlations and surfaces potential bottlenecks and issues.

Level 4- Federated Observability:
Federated observability enables the democratization of observability. Data is made available for consumers on demand. It ensures business agility at the lowest possible cost. Workflows, consumption models, cost management, and other factors lie at the forefront at this stage. It’s the stage where data reigns supreme and data consumers take the center stage leveraging data on-demand to solve a wide variety of security, operational, and business problems and make critical business decisions.

It assists organizations in fulfilling industry standards, functioning efficiently, and better-serving business goals.

Over the years, observability practices have evolved to rectify these difficulties, blending monitoring advances with a more comprehensive approach that provides deeper insights and a better understanding of what's going on across IT infrastructure.

The Observability Maturity Model breaks it down into 4 distinct levels (and beyond) in the development of observability. Let's have a thorough look at each stage.

A brief overview of the Observability Maturity Stages

A commonly accepted Observability Maturity Model has been proposed based on studies and discussions with businesses from various sectors.

The four stages of the current Observability Maturity Model include Monitoring, Observability, Active Observability, and Intelligent Observability.

Each stage of observability builds on the groundwork established in the preceding stage to provide further functionality for recording, tracking, and analyzing data with more efficiency.

With every other stage or level, the new features build deeper observability patterns at each stage, resulting in improved IT reliability and client satisfaction.

However, a key limitation of this model is that it needs to account for changing needs as organizations develop and change. As enterprises mature, they may require more data to move from one maturity level to the next; thus, gathering additional types of information to advance from one maturity stage to another isn't unheard of today.

What are the key elements of a good Observability Maturity Model?

The main elements of the ideal observability model comprise the already mentioned four levels. However, as data becomes more and more voluminous and as the standards of the industry rise, it's not going to take much longer for additional stages to be included.

The current model still focuses on the idea that Observability is a single-vendor concept. The next stage in maturity would be one in which a particular vendor platform isn't always sufficient, and enterprise maturity is going to be about how quickly you can adopt new solutions on demand.

Workflows, consumption models, cost management, and other factors lie at the forefront of the new model. It’s where data reigns supreme and data consumers take the center stage.

Let’s have an in-depth overview of each stage in the observability maturity model:

Stage 1: Monitoring (Is everything in working order?)

The IT environment is no stranger to the first stage of the Observability Maturity Model, Monitoring. Moreover, as reliable system operation becomes more important, the significance of monitoring grows as well.

Monitoring answers the simple question of “Whether the individual components are functioning as expected or not?” Monitoring is the process of analyzing a pre-determined set of numbers and failure situations. It monitors the component-level metrics including performance, capacity, and availability, and issues alerts if a monitored value is changed.

Events are important in that they are shifts in the IT ecosystem that require attention. Events usually represent significant happenings that call for involvement, despite the fact that they may just be instructive in nature. Moreover, Email, chat software, and mobile apps are just a few examples of the different ways that notifications or alerts can be created in reaction to occurrences.

Monitoring lets you know how your system is performing, whether any components are failing or breaking down, and what the status of each one is. It's a crucial first step that lays the groundwork for more advanced monitoring techniques.

In a nutshell, monitoring is about the following things:

> - Monitoring the general health of each component in an IT system.
> - Examining events and setting off alarms and notifications.
> - Notifying you that a problem occurred.

An administrator, for example, would install an agent on a server to track its usage. The data from the agent is collected by a management server and displayed via the IT monitoring system's user interface, usually as a graph of performance over time. If the device stops functioning properly, it sends out an alert to the admin; he or she can repair update, or replace it until it meets standard operating requirements.

Thus, Monitoring alerts you when an anomalous situation occurs, allowing you to get radical insights into the status and health of individual components. Therefore, it's a vital first step that sets the foundation for further observability development.

Stage 2: Observability (Why is it not working?)

Observability applies the same principles as monitoring to a much-advanced level, allowing you to discover new failure modes.

To make an analogy with Don Rumsfeld's response to a DoD news briefing question on February 12, 2002, Observability goes beyond what you know. It doesn't anticipate that you'll have a clue about the source of an effect seen in your application data. There needn't be even an event for observability to function. At its core, it allows you to identify and comprehend things about which you can't predict failure modes in advance.

You'll need a thorough understanding of what's going on with your system in order to figure out what happens when an alert appears. Observability usually delivers these insights by concentrating on three important categories of telemetry data: metrics, logs, and traces.

The following three foundations of observability are drawn from IT components such as microservices, apps, and databases to provide a systems-level view into the functioning of a system:

- **LOGS:** The term "log" refers to a file that records events, warnings, and errors as they happen within a software environment.

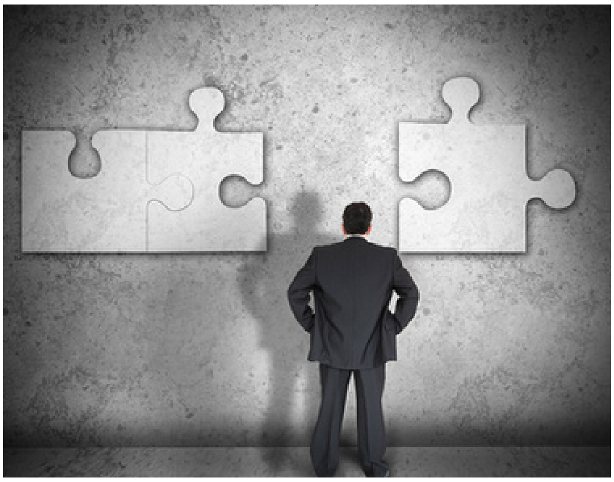

- The majority of logs include contextual information such as when an event took place and whom the user or endpoint was associated with it.

- **METRICS:** Metrics are numerical measurements that assist you in understanding the performance and condition of your services. The four golden signals comprise latency, traffic volume, error rate, and saturation.

- **TRACES:** A request's journey from beginning to conclusion is depicted in traces, which are in-depth representations of how data flows through an application. Traces aid in performance troubleshooting and occasionally provide insight into your application's performance at the code level.

Additionally, dashboards are frequently used to display metrics, logs, traces, events, and alarms so that developers can conveniently keep track of important actions.

Site reliability engineers and DevOps teams can generate auto-generated dashboards based on dashboard templates or create custom ones by scratch.

Furthermore, some exclusive monitoring systems offer combination panels that compile these different types of data under one roof, viz. in one system, and allow for deeper visibility.

Stage 3: Active Observability (What is the origin of the problem, and what are its consequences?)

Observability at stage 2 is good but not without its shortcomings. Data that is generated when implementing observability is extremely voluminous and often times its hard to segregate the useful from the redundant ones.

Dealing with Data siloes and volume with stage 2 observability soon becomes a pain in the back. In order to diagnose a problem, you might need to build arbitrary solutions that query various observability silos; creating these queries requires developers to have development skills, in-depth data structure knowledge, and a thorough understanding of system design.

Moreover, to facilitate the collection of metrics, logs, and traces, businesses have widely adopted models like OpenTelemetry and Prometheus.

They are quite useful for collecting data but when it comes to unifying silos or providing a better context for data, stage 2 observability simply falls short.

This is where Active observability comes in. You will need to contextualize events, logs, metrics, and traces from across the data silos in your infrastructure to discover how your observability data is connected.

The Apica Ascent platform combines active observability with load testing, streamlining cloud migrations, application management, and infrastructure issue resolution.
Apica's active observability solution offers enhanced capabilities to unify complex data sources through the operational data fabric, optimizing insights up to 95% faster, and reducing costs.

With a comprehensive suite of products, namely, OBSERVE (including synthetic monitoring and load testing), FLOW, and LAKE, Apica helps alleviate dashboarding friction, empowering customers to access and utilize data seamlessly with the powerful integration of active observability

What is Topology?

Modern environments are made up of a number of dynamic layers, including microservices, serverless applications, and network technology. So it becomes imperative to include an up-to-date topology in your observability stack.

Topology is the most significant factor affecting Active observability. The topology of your IT system is a representation of all the components at every layer—from network to application to storage—and how they all function together. Topology takes into account logical relationships, vicinity between members, and other linkages between components to provide clear visualizations and operationalized connection data.

Simply put, the purpose of topology is to group and cluster a range of data streams, revealing previously hidden interconnections. Additionally, Topological visualization gives you crucial information about how your business is impacted when something goes wrong and enables you to analyze network telemetry.

Adding topology is a good first step albeit insufficient in and of itself to provide Active observability. Time is the second most important factor here. The incorporation of topology and time at stage 3, allows you to discover the source and impact of any change or failure in different layers, data silos, teams, and technologies. Moreover, You can automate root cause analysis, business impact analysis, and alert correlation with stage 3 observability.

Stage 4: Intelligent Observability (How to predict anomalies and automate response?)

At stage 4 observability, AI/ML algorithms look for patterns signaling errors correlation and remediation workflows driven by AI. In other words, observability is intelligent at this stage.

AI and machine learning (ML) are used to analyze huge amounts of information in the context of monitoring and observability.

AI/ML algorithms look for changes in patterns that indicate impending warnings, alerts, and failures to help businesses detect when a service or component begins to deviate from normal behavior and address the situation before something breaks. This is where AIOps come into the picture. According to Gartner's Glossary, "AIOps combines big data and machine learning to automate IT operations processes, including event correlation, anomaly detection, and causality determination."

The next level of maturity in the AIOps architecture, known as Analytics for Operations (AO), focuses on feeding data into an AI framework.
The goal is to deliver more precise suggestions for resolving issues within this platform's stack and provide greater-than-human intelligence solutions by utilizing machine learning technology.

This layer also builds off previous levels' capabilities such as collecting and processing information, topology assembly, and data correlation - adding pattern recognition, anomaly detection, and other refined recommendations for remediation.

How do AIOps operate?

Elements of IT operations and application monitoring can be automated and made simple with AIOps platforms in the following ways:

Data collection and analysis: All logs and events produced by infrastructure and apps are collected and examined. Following analysis, AIOps systems highlight data that indicates a problem.

Pattern detection: Discovering patterns involves correlating and identifying correlations between various data elements using AIOps platforms.

Disruption: AIOps analyzes the underlying causes of both new and persistent problems so that enterprises can prevent their effects by taking proactive measures.

Collaboration: is facilitated and made easier by AIOps platforms' unified dashboards and sophisticated notification frameworks.

Automation: AIOps works to automate responses to issues and threats as much as possible so that issues and threats can be resolved quickly and easily.

A few of the key benefits attributed to stage 4 observability include the following:

- Deep insights into how the operations of the IT environment work utilizing AI/ML to collect and correlate useful information from vast amounts of data.
- Anomaly detection and predictions that identify problems before they have an impact on the business.
- Improved productivity and less work as teams concentrate on the most important events
- Increased accuracy of alert correlation, performance review, and intelligent root cause analysis.

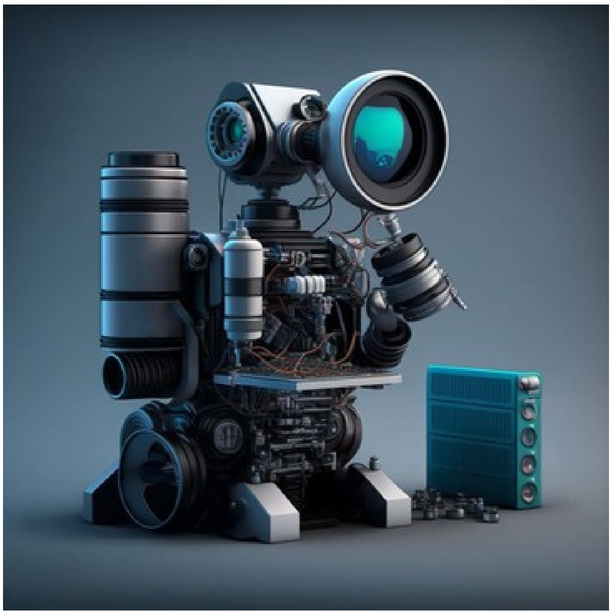

Additional Benefits

Improved efficiency: Automated data analysis and machine learning can help to identify issues more quickly and accurately, reducing the time and effort required to troubleshoot and fix problems.

Enhanced visibility: Intelligent observability provides a more comprehensive view of system performance and behavior, allowing organizations to better understand how their systems are functioning and identify potential issues before they become critical.

Better decision-making: By providing deeper insights into system performance, intelligent observability can help organizations to make more informed decisions about how to optimize and improve their systems.

Increased reliability: By identifying and addressing issues more quickly and accurately, intelligent observability can help organizations improve the reliability and stability of their systems.

So, Stage 4 in the observability maturity model is when your organization should experience increased IT operations that provide a superior consumer experience. Set up Intelligent observability to traverse silos and ingest data from across the environment in order to reach these objectives. The AI/ML models should evaluate all of the observed phenomena, metrics, logs, traces, changes, and topology we discussed previously based on their relationships over time.

Stage 5: Federated Observability (How to make it accessible and available to all?)

The next step in the observability data model deals with the idea of open observability. In other words, data availability for consumers with on-demand convenience.

Up until this level, observability is constrained by the viewpoint of a single-vendor paradigm. Now at stage 5, observability continues to support hybrid, multi-cloud architectures as they spread toward the edge and incorporate machine learning, microservices, containers, and other cutting-edge technologies. The differentiating point lies in how the line of accessibility is being redefined.

The objective of developing and incorporating the 5th stage into the model is the democratization of data. This brings about better workflows, conspicuous consumption models, and enhanced cost management practices besides a plethora of other elements.

It is pretty much evident that observability is essential for maintaining your entire digital ecosystem, which means it's no longer a distinguishing feature. It's an important competence that all businesses must master to keep their operations going. So, why does observability has to be exclusive? Why is it so exorbitant? And most importantly, why are the majority of organizations are still struggling to achieve complete observability?

With the new observability maturity model, it’s anticipated that the dangling backlogs of the earlier stages will be rectified. Stage 5 observability is imperative, in that it has become the need of the hour. With Web3 around the corner, almost all online data is going to be decentralized. To keep up with the unprecedented shift in volume and security, therefore, federated observability is bound to be embraced by the community.

How Apica is supporting Federated Observability

In today’s cloud ecosystem, achieving 100% observability is capturing lightning in a bottle. Identifying the appropriate monitoring approach to manage one's environment consistently is difficult for businesses.

Monitoring has been used for years by IT operations teams to learn more about the performance and availability of their systems. However, because of the multiple dynamics, dispersed, and modular IT environments that today's digital infrastructures and apps span, a more in-depth understanding of everything that occurs in these systems is required.

Observability provides a thorough understanding at every step of maturation by giving distinct capabilities. But as the challenges evolve, more compliant stages need to be included in the maturity model accordingly.

Your IT systems will be more dependable and durable as you advance in the maturity model. You'll be able to pinpoint the root of problems more quickly, appreciate how modifications and malfunctions impact business operations, and ultimately give customers a better experience.

With the help of technologies like AIOps and Machine learning, most organizations have achieved intelligent observability. However, there’s still an accessibility gap that’s only becoming wider with new challenges like data sprawl, overflowing machine data, and rising security concerns. Apica is committed to fill in the gap with its AI-powered infrastructure that’s also in compliance with the notion of federated observability.