

Understanding and Achieving Observability - A Primer

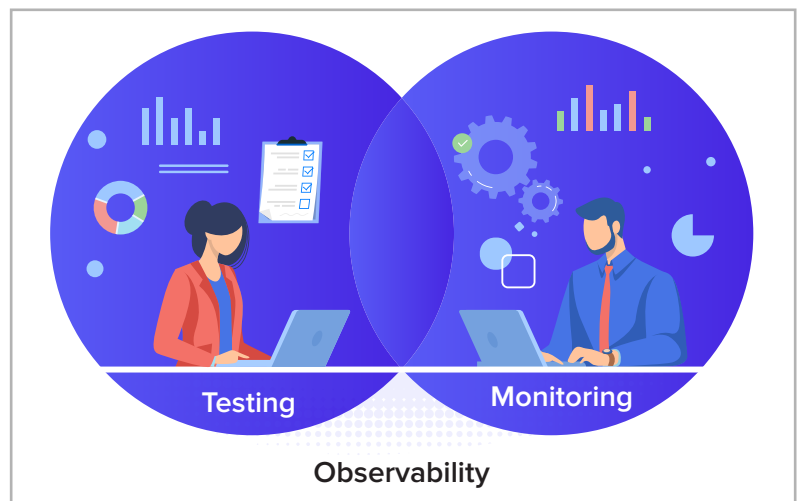


Synopsis

Observability is the talk of the town in the DevOps community these days. While IT veterans argue what the definition of observability is, it's important to note that observability does not exclude monitoring. In fact, observability is a superset of testing and monitoring enabling you to have a holistic view of your application's behavior and health at any given point of time. This white paper aims to discuss the definition of observability, its evolution, what it means from person to person, and how to achieve it.

Introduction

As the role of DevOps evolves from being reactive to operational issues, to being able to proactively identify issues, it's critical to prepare systems for effective monitoring and correlate data from all available sources. Here we take a look at yet another evolving trend on making systems observable.



What is observability?

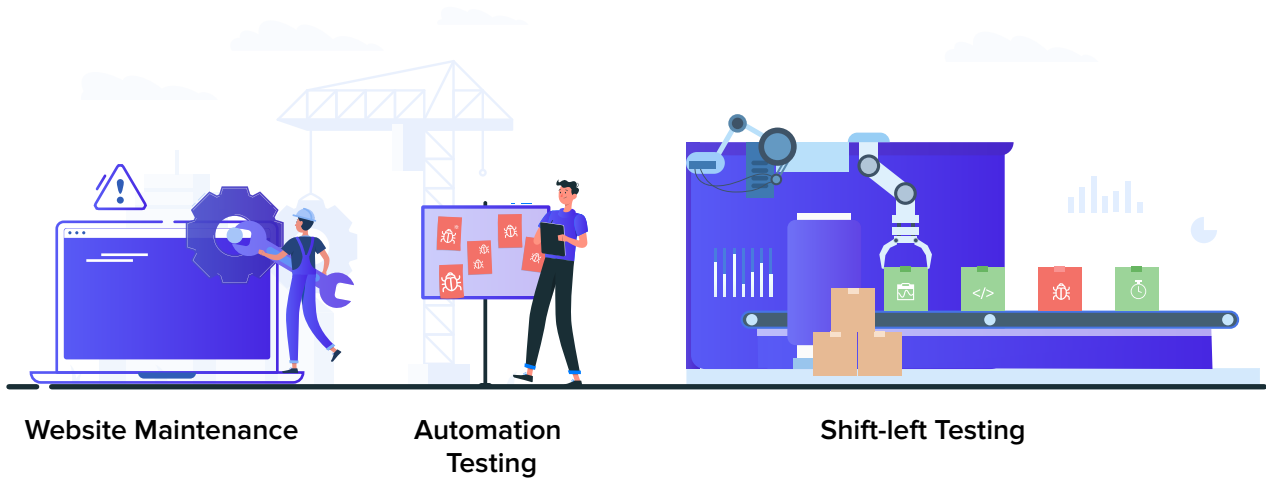
It's hard to define an evolving buzzword as each segment of the industry associates differently to it. For the purpose of this article, we define observability as:

“A state of readiness of the DevOps team to instantaneously identify unexpected application behavior and pinpoint its root cause.”

Shift-left Testing and Observability

The utopian dream of any DevOps team is to have no unknown issues in their production environment. When an issue is identified, DevOps teams try to understand why it happened and include it as a test case in the testing process. This is commonly known as shift-left testing. The shift-left testing is important for DevOps to stay ahead in the competitive landscape.

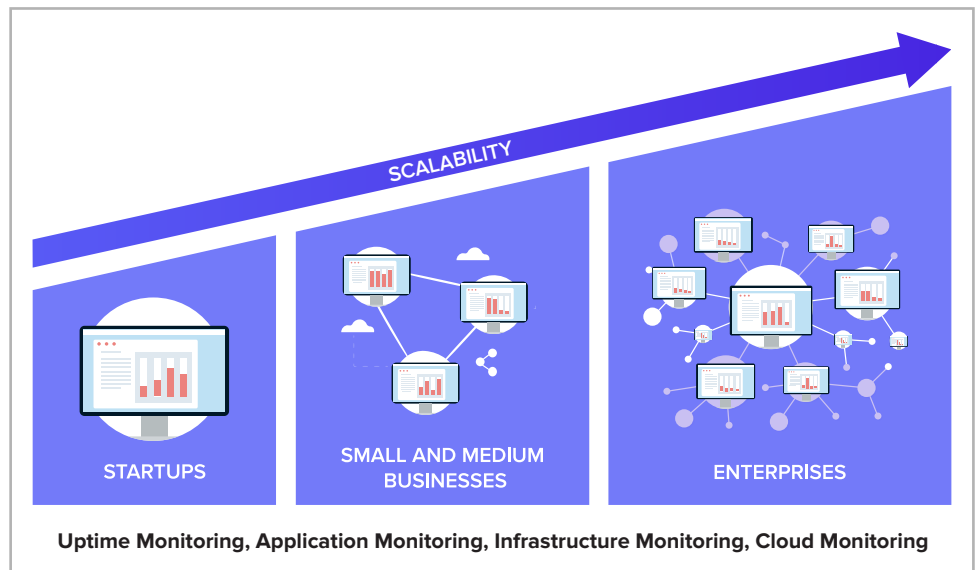
With the growth in monitoring and automation tools, DevOps teams are able to "shift-left" their quality assurance (QA) and testing process. By making systems more observable, DevOps can easily understand how to quickly shift-left any new problems.



Setting Goals for Observability

The monitoring needs of an application/service largely depend on its scale and size. A small-scale application may not need much monitoring, but as the scale and complexity of applications/services grows, so does its monitoring needs. This is why the scope of monitoring and observability depends on your organization's size and the scale of its services.

For instance, a web hosting provider may only need website and infrastructure monitoring to notify its customers when its service becomes unresponsive, while the customers themselves may need log and the application performance monitoring to keep up with their SLAs. On the other hand, a more sophisticated payment gateway service that handles millions of transactions a day would need continuous



monitoring of all transactions, logs, and business metrics, with smart alerting and response systems that can perform automated actions during failures and generate reports on the root cause analysis for outages.

Monitoring goals vary according to an organization's specific needs, which means the goal of observability depends on what you want to monitor; it may vary from person to person within your organization, and will differ from other organizations' goals.

Achieving Observability: Pillars and Pitfalls

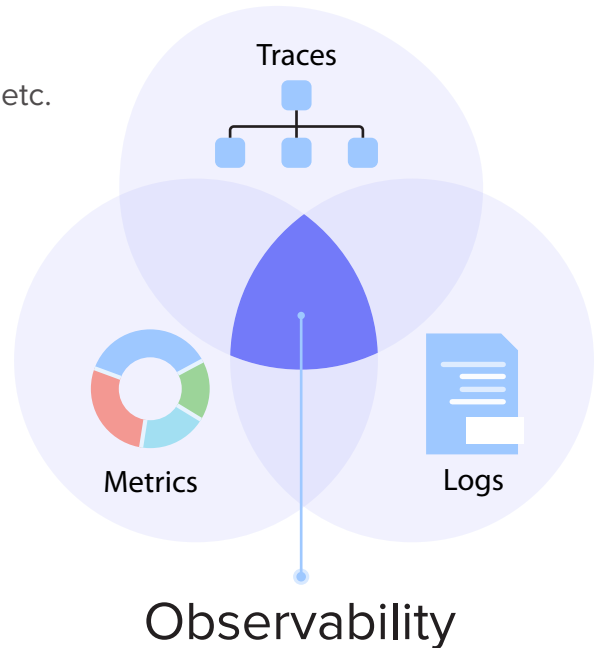
Once you've defined your goals of monitoring and observability, you have to focus on refining your application life cycle process to achieve these goals. The three main areas of focus, commonly known as the three pillars of observability, are metrics, logs, and traces.

1. Metrics: There are industry standard metrics like Apdex, average response time, 95th percentile lines, etc. to get you started. But you can go further by adding business-specific metrics like scheduled jobs, third-party resource utilizations, key component usage, etc.

2. Logs: Application and server logs can provide request-specific details that cannot be captured via standardized monitoring solutions. While app server logs provide request-specific details, such as user agent, features accessed by the users, and IP addresses, developers can add contextual information about the state of the system. Analyzing the logs tells the story of system behavior over time.

3. Traces: Debugging slow transactions requires insights into time-consuming methods and external services.

Traces provide a snapshot of the entire workflow of a transaction, making it easier to identify parts of the code that require optimization and fine-tuning.

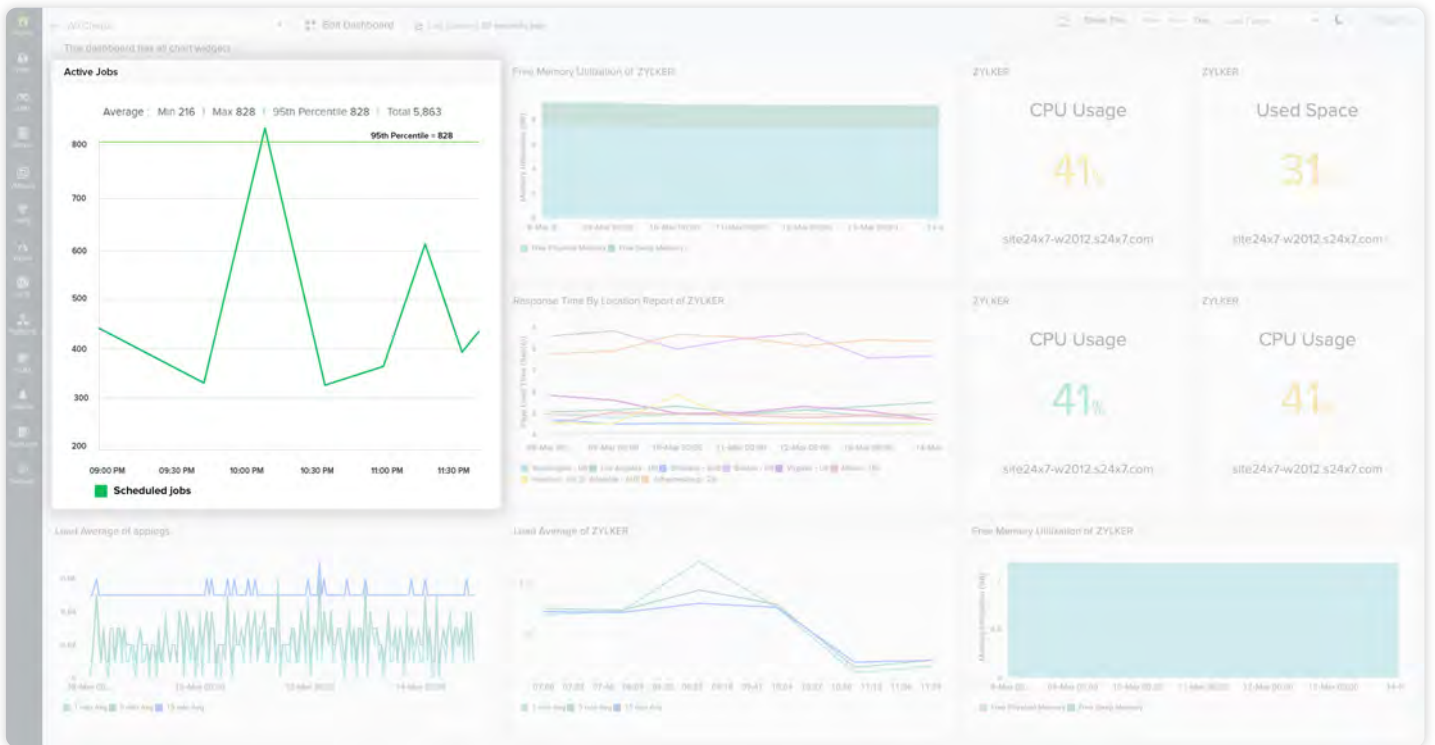


Achieving Observability: Guidelines

Once the goals of observability are established, it's time to align your monitoring solutions with them. Following these guidelines, you'll be able to get the most out of your monitoring solution and tailor-fit it for your needs.

1. Use custom metrics: Monitoring tools have done a great job to get you started in troubleshooting by capturing numerous industry standard metrics. However, they can only provide you an outsider's perspective on your service. Defining and monitoring application-specific metrics may increase the development workload, but it's definitely worth the trouble since it can provide deep context into the different states of your system to help you understand what factors need to be tuned to achieve optimal performance.

Site24x7 manages a lot of scheduled jobs to run customer-specific workloads. By monitoring the active jobs, we can determine how to scale the servers and improve resource utilization. Since the jobs are running as separate threads spawned across multiple servers, monitoring them is only possible by defining custom metrics.

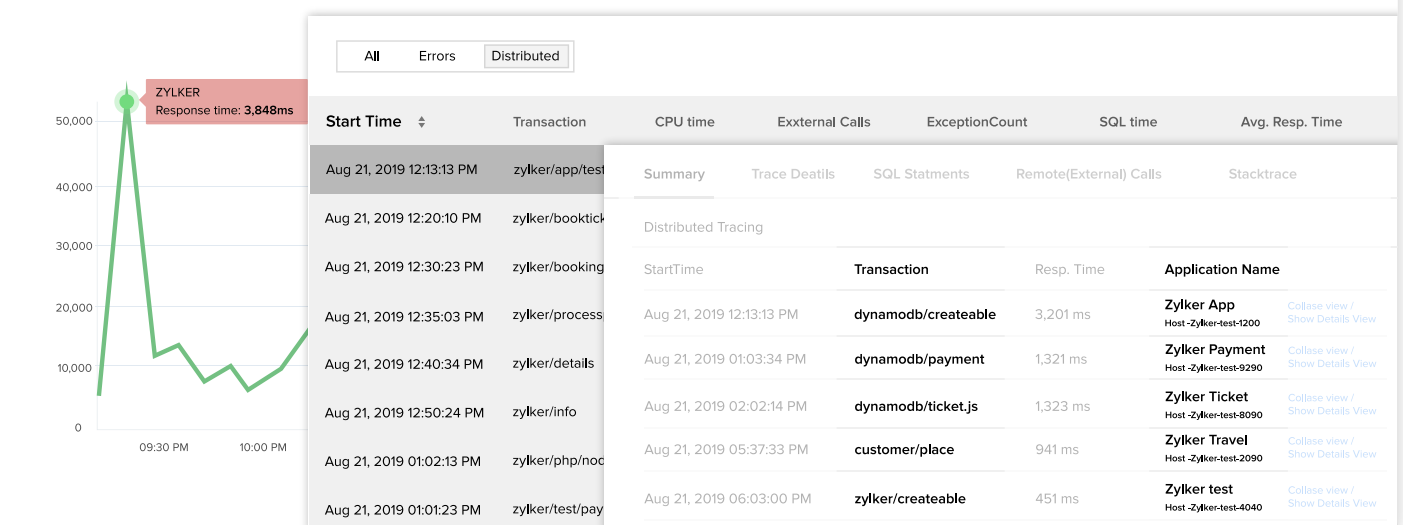


2. Analyze traces: Monitoring slow application programming interfaces (APIs) is a great way to increase observability of a system. Dig deep through traces to discover bottlenecks in the performance of your application/service. If a third-party is slowing you down, find ways to replace or avoid the third-party service. If the database is slowing you down, consider upgrading or de-normalizing your database. If you have a distributed system, use a monitoring solution that can provide distributed traces. With distributed traces, you can not only optimize individual applications, but rework the communication flows and enhance the entire network.

Site24x7 offers distributed tracing, allowing you to monitor code flows across application boundaries.

Distributed Traces

Response time of zylker/app/test.svc

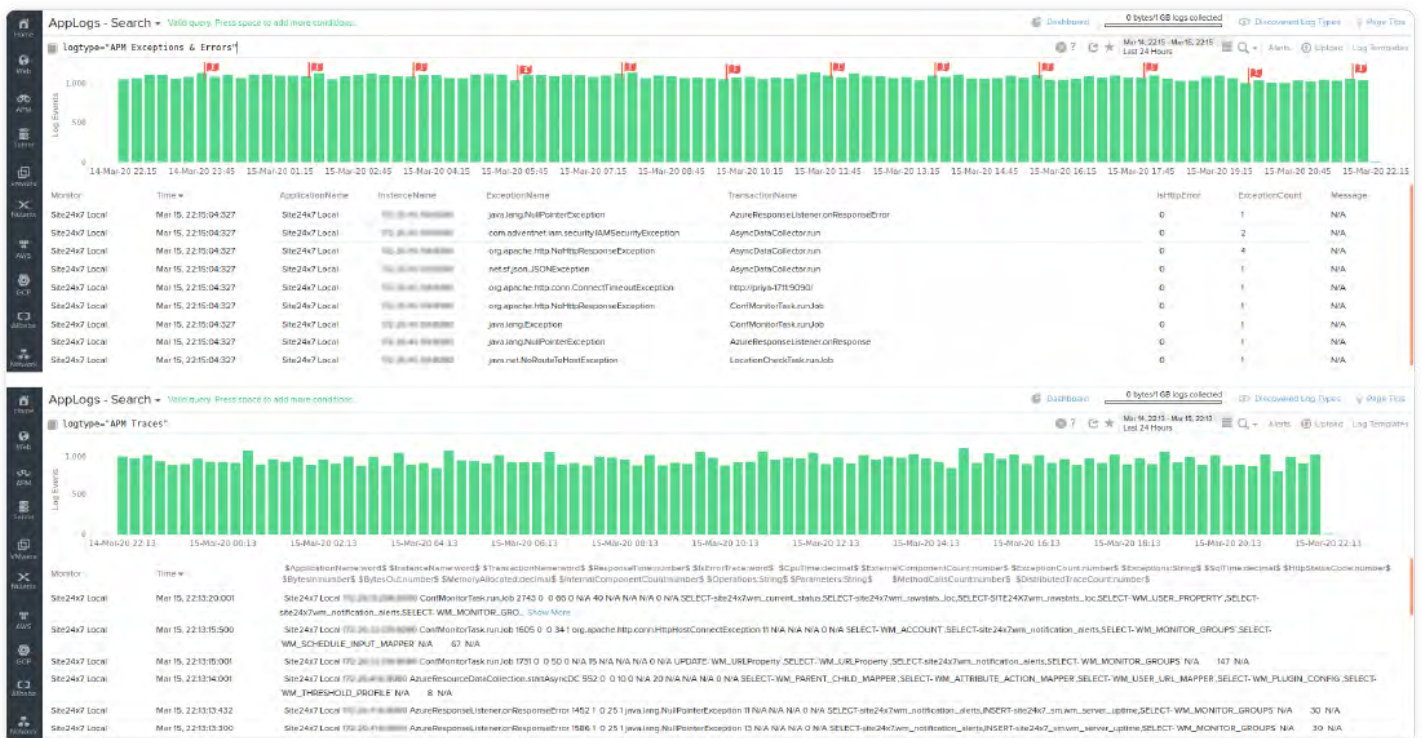


3. Structure logs for data mining: While metrics and traces alone can give a holistic picture of your application's day-to-day operational behavior, they are limited to monitoring internal factors. Your application performance may be affected by external factors like user behaviors, which cannot be captured by metrics and traces.

This is where logging can provide context. Is a single rogue agent choking your bandwidth? Is application performance only taking a hit from certain geographies and browsers? Is the database constantly bombarded with the same query, the results of which can be cached? These questions can only be answered by analyzing logs. Log monitoring makes answering these questions a breeze. However, without the context of traces and custom metrics, the answers to these questions can very well lead your performance optimizations the wrong way.

If you structure your logs so that they can be queried, associating them with traces and other metrics becomes more natural and relevant. Restructure your logs to include custom metrics and use traces to identify where logging is more relevant. Once you bring the context of custom metrics and traces into your logs, you will have a holistic view of your application's state. You can then feed this data to tools that can perform further analysis.

In Site24x7, all service logs are pushed to our central log server from which we query and monitor infrastructure events. Site24x7's Log Management console provides querying functionality along with trend analysis, which makes it easier to correlate and arrive at root cause analysis (RCA).



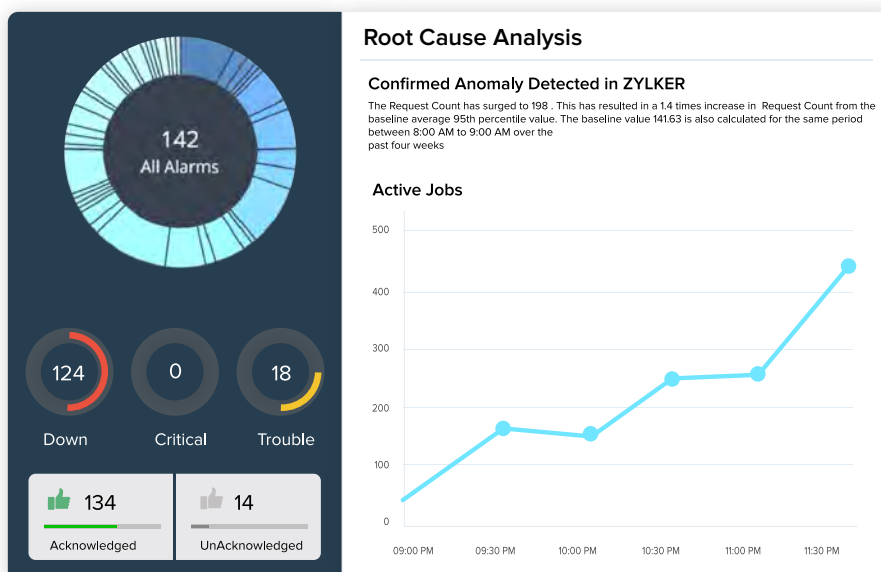
If you structure your logs so that they can be queried, associating them with traces and other metrics becomes more natural and relevant. Restructure your logs to include custom metrics and use traces to identify where logging is more relevant. Once you bring the context of custom metrics and traces into your logs, you will have a holistic view of your application's state. You can then feed this data to tools that can perform further analysis.

4. Perform trend analysis: Having a lot of structured data is only the starting point of performance tuning. You also need complete visibility on how each change affects the overall performance of your system. To see this, you need a solution that can store historic data and perform trend analysis, preferably with some ML/AI built-in to detect anomalies and alert you whenever the system does something unexpected. After gathering enough data-points and identifying the relations between them, you can even begin to automate some decisions.

5. Set up alerting and automation: An ideal alert system should be intelligent enough to distinguish patterns of application behavior under normal and extreme circumstances. Getting bombarded with alerts that end up getting ignored is just as good as having no alert system at all

For setting up intelligent alerts, many tools integrate ML/AI to study the application behavior over time and detect anomalies. However, the best alerting system can only be built by carefully configuring critical parameters based on experience. For instance, if your disks are getting full frequently, you can set up alerting when disk utilization reaches 80 percent, and set up automated scripts to run when thresholds are breached, like archiving old data or clearing up old logs, cached files, etc. Understanding the system well enough to make these informed decisions comes from experience. However, with reports from monitoring solutions, these decisions can be made much more quickly, like archiving old data or clearing up old logs, cached files, etc.

Site24x7 has automated actions to not only mitigate a failure, like retrying failed jobs, but also to collect diagnostic information during abnormal circumstances, like memory dumps when usage is high. These actions make the job of DevOps teams easier, since they can attend to only critical events, and let the automation handle the rest.



Conclusion

You get to decide what level of observability you want in your system. Use your resources, instrument your metrics, and set up proactive alerting to correlate information across various silos to get the best information possible. In short, observability is neither a tool nor a mechanism, but more of a set of practices/culture that helps you to align your goals.

About the Author

Ranjani Subramanian works as a Product Marketer at Site24x7. She is committed to generate meaningful and informative content that educates the readers. Her areas of interest revolves around application performance monitoring, digital experience management and end-user experience. While not at work, you can always find her talking nonchalantly about movies and Netflix.



Ranjani PS
Product Marketer

About Site24x7

Site24x7 is a full stack monitoring solution that empowers IT operations and DevOps with AI-powered performance monitoring and cloud spend optimization. Its broad capabilities help quickly troubleshoot problems with end-user experience, applications, servers, public clouds, and network infrastructure. Site24x7 is a cloud offering from Zoho Corporation, which has offices worldwide, including the Netherlands, United States, India, Singapore, Japan, and China. For more information about Site24x7, please visit <http://www.site24x7.com/>