

Metrics and Observability: A Field Guide

A Field Guide



BY JOANNA WALLACE

3 THE INSIGHTS IN THIS EBOOK

<PART 1>

4 THE CASE FOR METRICS

5 WHY DO YOU NEED OBSERVABLE METRICS?

6 THE ROLE THAT METRICS PLAY

6 Application metrics indicate the health and load of your application

7 Product quality metrics show characteristics of your product

7 Project metrics describe project and execution characteristics.

8 Marketing metrics show user interactions.

9 WHAT IS OBSERVABILITY, AND HOW DOES IT AFFECT ME?

10 HOW TO ACHIEVE AN OBSERVABLE SOLUTION

<PART 2>

12 THE MOST COMMON MISTAKES WITH METRICS

13 Graphing averages, medians, and percentiles

14 Stacked versus overlapping bar graphs

16 PRACTICAL MISTAKES TO AVOID

16 Creating monolithic solutions

17 Use the best graphic for each data set

<PART 3>

20 THE SOLUTIONS ON THE MARKET TODAY

22 REGULATORY COMPLIANCE

22 ALERTING AND ALERT TYPES

23 What is Alert Fatigue, and why should it be avoided?

23 Static versus Dynamic Alerting

24 Specialized Event Alerting

25 DEFINING DYNAMIC METRICS

25 Types of Metrics Available

26 VISUALIZATION OPTIONS

26 Proprietary Solutions include Custom Dashboards

27 Plugins are also available for external visualization tools

27 PRICING

28 The listed cost of proprietary or the hidden cost of open-source?

28 Cost of Data Quantity and Internal Processing

30 WRAP-UP

THE INSIGHTS IN THIS EBOOK

An observable system is one where you know exactly how well it is performing at any given time. Observability is gained by collecting and analyzing different data sets; using them gives insights into your system's health. In this eBook, we will take a deeper dive into one [pillar of observability: metrics](#). We will explain how metrics can help improve the usability of your system, explain common mistakes made while creating metrics, and dive into existing observability solutions to help you choose the right tools for your software solution.

<PART 1>

THE CASE FOR METRICS

Before deciding how to implement metrics, you need to know why they are critical for your product. When metrics are available, they will change how you approach problem-solving by providing qualitative evidence to support or dissuade specific decisions.



WHY DO YOU NEED OBSERVABLE METRICS?

When your service is down, your bottom line takes a hit. The whole goal of the tech industry is to provide the ideal solutions that people need. Of course, in practice, there will always be problems.

In [February 2017](#), a human error caused an AWS outage, costing them and the S&P 500 companies using the platform an estimated USD 150 million in lost revenue during the outage. In June 2021, [Fastly](#), a content delivery network, experienced a bug-caused outage that resulted in [their customers](#), such as PayPal, Spotify, eBay, Amazon, and thousands more, losing revenue, productivity, and IT recovery.

These vendor outages are just one potential cause for your solution to be down - meaning your customers can not view, use, or purchase from your interface. This is where observability comes into play. Observability means knowing when your solution is down or at risk to go down (completely or partially), understanding why it is down, and knowing the steps to correct it. One study on [microservices adoption](#) claimed that 27% of organizations listed *monitoring and observability* as one of their biggest challenges. A further 16% of respondents said debugging was the biggest challenge. The observability industry has been rising and meeting the needs of its cloud-solution clients.

OBSERVABILITY CHALLENGES

Flood of data from a constantly changing distributed system

Fault detection lies with a select few team members with deep knowledge of the system

Fixes are reactionary to user complaints instead of proactive

Niedermeier et al., 2019

THE ROLE THAT METRICS PLAY

Metrics are only one piece of your observability stack, working alongside other tools like [logs](#), [traces](#), and [security](#). Well-designed [metrics](#) can give any stakeholder or developer a fast, easily interpreted insight into what is not working for their customers. A variety of metric types can work alone or in concert with other metrics and resources to meet company goals. When metrics measure your system's quantifiable aspects and are reproducible and objective, they can be an invaluable asset that influences your roadmap.

Developers can group metrics into several categories. We will consider application, product quality, project (or resource), and marketing metrics. Let's look at examples of each metric type to show why it is useful.

Application metrics indicate the health and load of your application

Application metrics look at how your host-level software is functioning. Specific metrics you will need to measure will largely depend on what your platform does and its dependencies. Application metrics tend to be user-defined to accommodate the customizations needed to understand your particular solution. However, there are some common metrics that most

or all software solutions will find useful. Common application metrics include:

- **Error and success rates** - How often are your functions being used, and how often do they fail?
- **Response latency** - How long do your functions take to respond to client calls?
- **Memory** - How much memory do your functions use?
- **CPU usage** - How much compute power do your functions use?

Product quality metrics show characteristics of your product

Product quality metrics are measures of software features and how they are functioning in the wild. If you have a software offering, you probably have competitors in the same space. To keep your customers and attract new ones, you need to have a quality offering that does what your users want. Having poor quality software can frustrate users and send them to your competitor. These metrics help you ensure you are keeping up the standards of your offering and can alert you when you are not, so teams can fix problems quickly before any user complaints arrive. Product quality metrics include:

- **Reliability** - How well does your software work, and how long will the system run smoothly without crashing?
- **Maintainability** - How difficult is it to maintain your software? Difficult to maintain software affects how quickly your team can react to problems.
- **Availability** - How often is your application available to users? The duration of availability is commonly known as uptime.

Project metrics describe project and execution characteristics

Project metrics indicate if your team can deliver and maintain software

quickly, efficiently, and reliably. These metrics typically use other gathered software metrics and statistics from your project management tools to calculate valuable data. Project metrics include:

- **Time to restore service** - How long does it take to go from an incident or error affecting user experience to a completed fix?
- **Change failure rate** - What percentage of deployments or changes released to users result in service degradation requiring remediation?
- **Lead time for changes** - How long does it take for completed code to be deployed and running in your production environment?

Project metrics help teams understand how they can better meet delivery targets and deliver more reliable software. Tracking these metrics over time can show if a change in your project management system has a beneficial or detrimental effect. Such measurements led [one report](#) to conclude that elite teams who could meet the delivery and reliability goals are 3.7x more likely to use some continuous testing methodology. These insights can lead your team to integrate specifically chosen tools that enable them to succeed.

Marketing metrics show user interactions

Marketing and SEO metrics help software and marketing teams understand how their features and upgrades affect users. They allow teams to measure the success rates of different campaigns, the usability of their sites, and whether or not targets are met. According to [a 2018 survey](#), 89% of leading teams use marketing metrics to track their products. Marketing metrics include:

- **Bounce Rate** - How many users hit your webpage but take no action once there?
- **Page Depth** - How many pages on your website do users visit during a session.
- **Session Duration** - How long do users spend interacting with your website?

- **Returning Visitors** - How many of the visitors on your website have visited before?
- **Conversion Rate** - How many of the visitors on your website take the actions you require of them?

Marketing metrics, especially conversion rate, tend to show the overall success of your website. While these metrics alone show how your website is doing, it is also critical to look at these in the context of other processes, product, and project metrics. Looking at metrics together allows teams to answer questions like *why the bounce rate to our website increased by 40% last Wednesday?* Other metrics alongside the marketing results can show teams that it may have been because of a latency issue whose duration coincided with the bounce rate increase.

WHAT IS OBSERVABILITY, AND HOW DOES IT AFFECT ME?

Cloud infrastructure production and maintenance rely on a constantly-changing knowledge of your infrastructure and how your users are interacting with it. Usage ebbs and flows, data is processed and changed, vendors provide upgrades, and physical assets decay. With all of these, the status of your system will change. Observability means you can quickly see how these factors affect your offering and adjust as needed with confidence, knowing your actions will be beneficial. More importantly, your team should be able to answer *why the changes are occurring* and adjust the system as needed.

Monitoring tools are abundant, visualizing surface-level data to show what is occurring in your system. On the other hand, observability tools have a slightly different goal - to give insights into why behaviors are happening in your system. Observability reduces the engineering time needed to isolate and fix problems, sometimes even [finding issues that wouldn't otherwise be detected](#).

When you hear about observability, you generally hear about logs, traces, and metrics. Each brings its advantages and information about your infrastructure, and when used together, you can obtain even more insights. Metrics are the observability tool best suited for statistical analysis, understanding overall system health, and alerting.

COMPLEXITY

90%

of surveyed IT practitioners agree that highly distributed applications create monitoring challenges an order of magnitude greater than other applications

80%

of surveyed IT practitioners agree legacy monitoring tools aren't adequate for monitoring modern cloud solutions

**VMWare State of Observability 2021*

HOW TO ACHIEVE AN OBSERVABLE SOLUTION

We've talked about the data you can collect to monitor your system. Monitoring means you might watch some specific metrics to see how your infrastructure is working. Observability looks at many, possibly changing, metrics in unison to analyze in context what is happening with your infrastructure.

Let's say you were one of Fastly's clients during their outage in June 2021. Fastly had a Content Delivery Network (CDN) outage that lasted up to 24 hours for some customers. With traditional monitoring, you would have some static metrics looking at the health of your infrastructure, which would have appeared normal with this type of outage. Your team may have also had monitoring on databases, compute functions, data queues, or API endpoints, which would have appeared fine.

An observability solution would go further than just monitoring, it would have detected the meager traffic to your site and notified your team of the issue. Observability solutions can use statistical analysis or machine learning to detect unusual changes in your data usage across your platform and notify your team to interject and solve the problem early.

To achieve observability, teams must either know which metrics are needed to solve problems before they happen or automatically use a tool that gives insights. Solutions [with built-in machine-learning alerting](#) can provide these insights to help your team solve problems instead of spending their time trying to find them.

<PART 1>

THE MOST COMMON MISTAKES WITH METRICS

Theoretical mistakes to avoid

When creating and visualizing metrics, a common issue is that teams will misrepresent or obfuscate their data.

We will discuss common visualization and calculation issues that can block teams from seeing bugs.


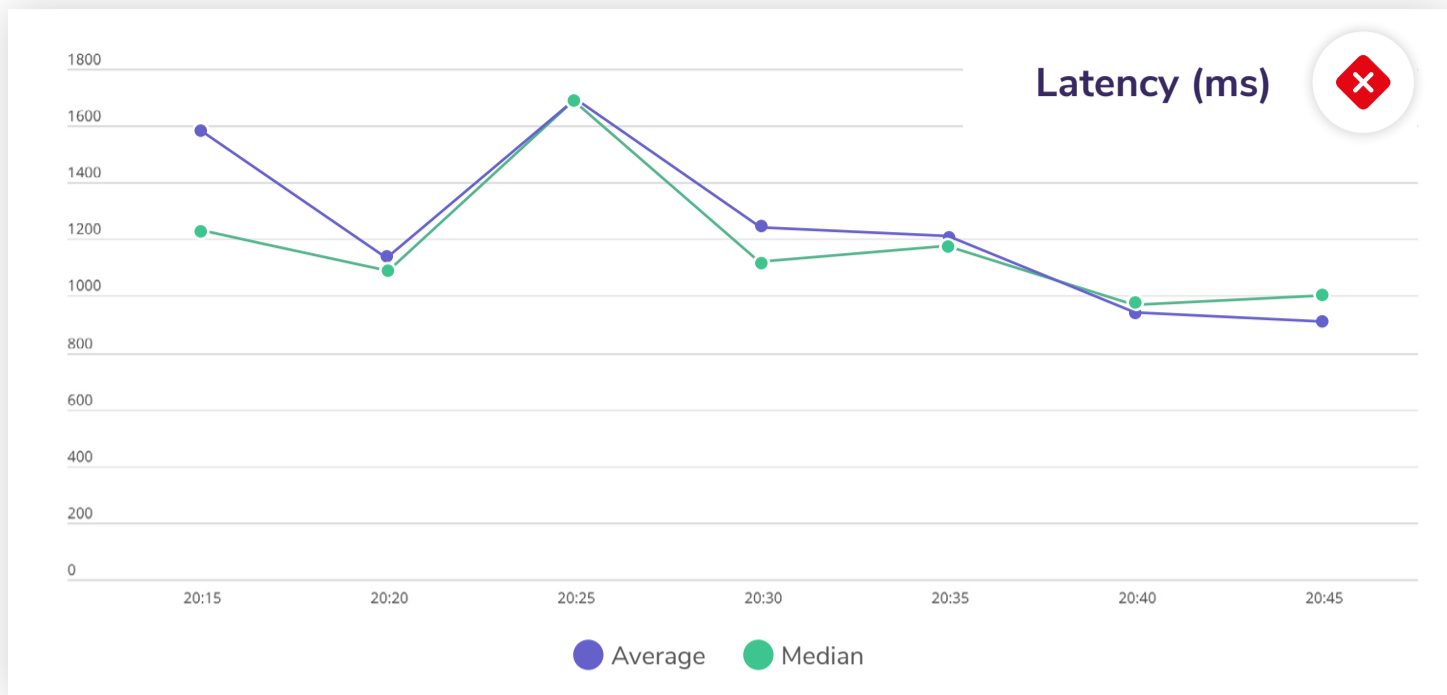


Graphing averages, medians, and percentiles

Almost 70% of consumers [claim](#) that the speed at which a webpage loads affects whether or not they will make a purchase. So, your engineering team should watch latency and make sure the numbers stay below a certain threshold. But, depending on how the latency values are graphed, they might not get a complete picture.

Graphing the average and the median can be helpful in knowing what the average user is experiencing on your site. But, by doing this, you will not see the worst latency times and will not know if those are less than your SLO. Let's assume your SLO says a particular page should load in under 2 seconds.

Below, the graph on the top shows the average and median latency for a page to load completely. They look like reasonable times and remain relatively consistent, with users, on average not experiencing latency of more than 1.7 seconds. However, the graph on the bottom adds in the 99th percentile calculation for the same data set. This shows a very different picture - that 1% of your users experience latency of over 4.5 seconds! If your team only sees the top graph, they may not take steps to correct anything since it appears the SLO is met. The below graphs will spur them into action to find why the latency is so high and correct it.




Stacked versus overlapping bar graphs


Stacked graphs are helpful when you need to look at the sum of different values and see the breakdown of those values. For example, if you are tracking conversion rates across various product lines, a stacked bar chart can show you the total conversion rate and each discrete product's conversion rate.

The stacked bar graph below shows the number of conversions for four products over a short period. Even with just a few discrete bars, you

can see that comparing the bars themselves is tricky. These graphs help compare a small number of values without knowing the exact values. They also clearly show the sum value and can indicate exceptionally high or low performers but are less effective when the data values are close.



If your goal is to see the difference between the performance of the different products, an overlapping chart is more effective. These charts overlap lines rather than stacking them, so it is easier to compare bar values at any given time. Similar to the stacked graph, it does have limitations in that this graph type can hide values that are the same. This may be acceptable if comparing a small number of values, but more extensive data sets can have hidden values with this chart type.



PRACTICAL MISTAKES TO AVOID

Creating monolithic solutions

When you attempt to see everything all at once, you risk displaying too much information and missing necessary data. Error metrics are an example of a critical metric to track. There are different ways developers may choose to visualize the error metrics, and some are more useful than others. When visualizing a metric, consider that the team using this data should be able to tell:

1. When did the failure occur?
2. How long did the failure last (or is it still failing)?
3. What service(s) failed?
4. How many users are affected by this failure?
5. Where can I look next to find out why it failed?

Logging the error metric on a simple line chart would be a simple visualization. Plotting an error count over time would satisfy the first listed requirement and show when failures occurred. But, how should you define the error metric graphed? There are some options for what could be done, and some are much better than others.

A monolithic solution would be to graph the total error count over time. While this would still satisfy the first requirement, it would not show what failed or indicate the next step towards a solution.

The next best solution would be to graph many lines on the same graph, each line showing the errors of a single function over time. This may be useful to see how errors in different functions relate to each other, but it can also be challenging to read. For example, if one function has five errors and another has 500 errors, the scale of the graph will need to be so large that the more minor error may not be seen. From the graphs

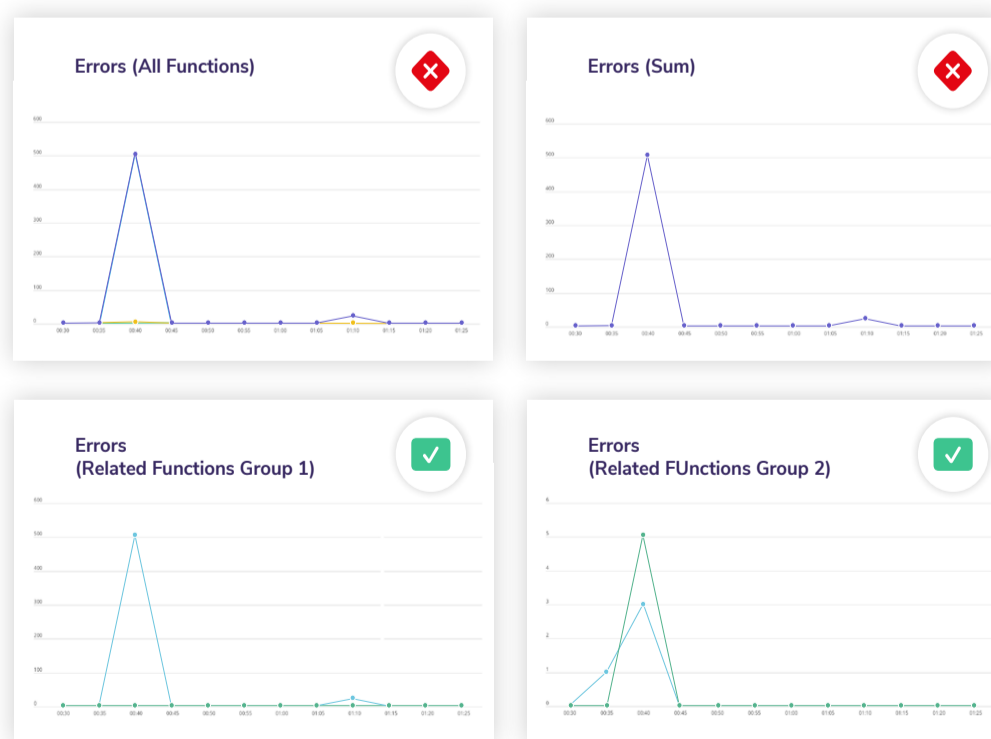
below, you can see the *All functions* graph was hiding some errors that are more easily spotted in the lower two examples.

The best solution would be to graph one or just a few errors on different graphs that appear on the same dashboard. Errors in each function or group should be easily seen if they run at the same scale. DevOps teams will also know from which graph shows error spikes, where in the system the error occurred, and where to go next to begin troubleshooting.


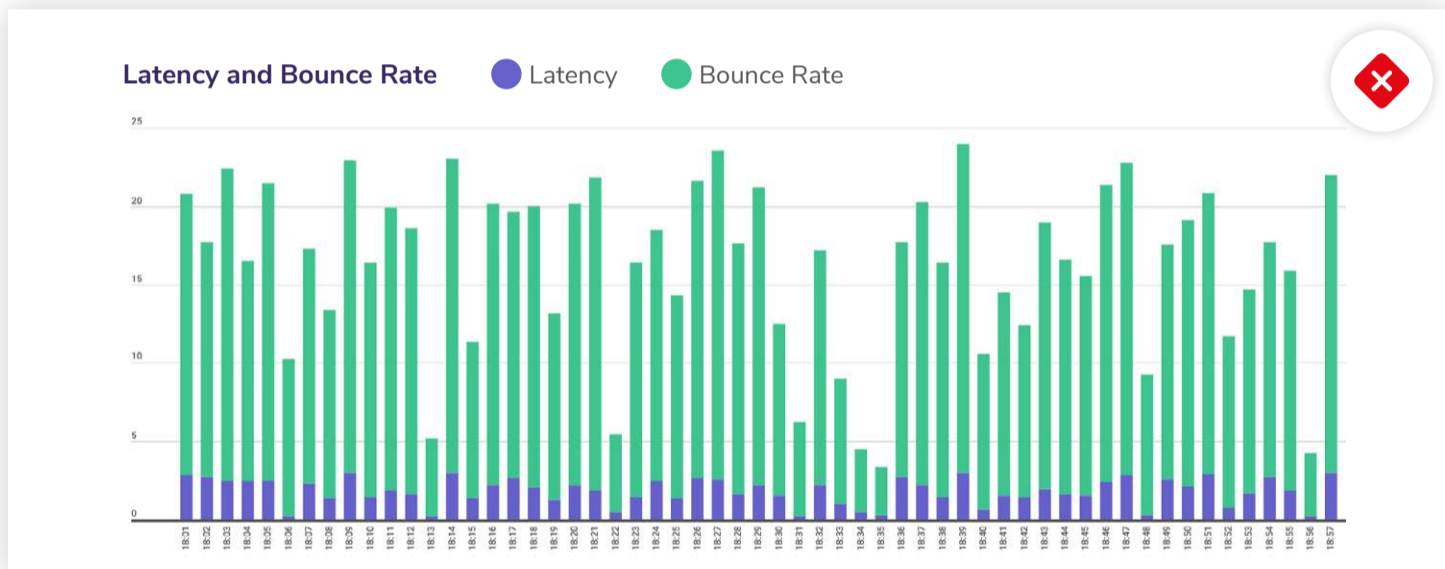

Use the best graphic for each data set

Each metric logged into your observability tool should be visualized to best suit it. In some cases, certain graph types cannot be used to display the data logically. In all cases, there is the best solution for how to show your data.

You might want to compare your known user bounce rate to the latency of loading your webpage. Your goal is to see if there is a correlation between latency and bounce rate. Which visualization will best represent this data? The three graphs below show different ways to represent this.




The first, a double-line graph, does not easily indicate if there is a relationship or not between these two values. Further, a single axis shows two different units of information (seconds and bounce rate). The second graph is a stacked bar graph. These graphs are meant to show both parts of each bar and the sum. Again, since these two values do not have the same units, the sum does not indicate anything and only hides the answer to our question. The last graph uses a line to show bounce rate and bars to show the 99th percentile latency over the same time segment. The graph shows two separate axes allowing the data to overlap and easily see if there are trends between the two data sets. Some of the highest bounce rates occur when the latency is highest, so this graph shows a correlation between the user and webpage behaviors.



<PART 3>

THE SOLUTIONS ON THE MARKET TODAY

There are many solutions on the market that provide some level of metric collection, creation, and alerting. The number of offerings to support debugging those solutions has increased with the popularization of microservice solutions. Here we will discuss some open-source and proprietary offerings on the market and how to decide which is suitable for your needs.



| | | CORALOGIX | PROMETHEUS | OPEN TELEMETRY | DYNATRACE | SUMOLOGIC | DATADOG |
|--|---------------|---|--|--|--|--|--|
| REGULATORY COMPLIANCE | GDPR | ✓ | None | Open-source solutions are managed either by users or by a cloud-managed service (like AWS). Regulatory compliance is not provided by the software inherently, but cloud providers may produce compliance | ✗ | ✗ | ✓ |
| | CCPA | ✓ | | | ✗ | ✗ | ✓ |
| | HIPAA | ✓ | | | ✗ | ✓ | ✓ |
| | ISO/IEC 27001 | ✓ | | | ✓ | ✓ | ✓ |
| | ISO/IEC 27701 | ✓ | | | ✗ | ✗ | ✗ |
| | PCI | ✓ | | | ✓ | ✗ | ✗ |
| | SOC 2 Type II | ✓ | | | ✓ | ✓ | ✓ |
| ALERT TYPES SUPPORTED | | Static (real-time) Dynamic (machine-learning) | Static (real-time) | None | Static (real-time) Dynamic (machine-learning) | Static (real-time) Static (scheduled) Dynamic (user input) | Static (real-time) Dynamic (user input) |
| DYNAMIC METRIC DEFINITION SUPPORTED | | ✓ | ✗ | ✗ | ✓ | ✓ | ✓ |
| VISUALIZATION OPTIONS | | <ul style="list-style-type: none"> ↳ Hosted Grafana ↳ Grafana Plugin ↳ Tableau Plugin ↳ Kibana ↳ Heroku ↳ Coralogix Widgets Dashboard | <ul style="list-style-type: none"> ↳ Grafana Plugin ↳ Expression Browser ↳ Console Templates | None | Dynatrace Widgets Dashboard | Sumo Logic Widgets Dashboard | Datadog Widgets Dashboard |
| PRICING | | Pay per GB for monitoring data, frequently searched data, and compliance data. Adjust your costs by configuring storage types | Free, open-source software, but users pay for the cost of the server and the engineering cost to set up and manage the servers. Managed servers are available from cloud providers like AWS at an extra cost | | Pay per GB, per host, and per internal data unit used. Features distributed across six discrete packages | Pay per GB, per host, per digital processing unit, and others. Features distributed across eight discrete packages | Pay per GB, per host, per test run, per log event, etc. Features distributed across fourteen discrete packages |

REGULATORY COMPLIANCE

Regulations associated with cloud and data storage are numerous. While many of us do not want to think about regulations often written in obfuscating language and expensive to implement, these are present to protect against serious data breaches. Typically, different regulations will apply based on your location, the sensitivity of your data, and your industry.

To choose the best solution for your platform, you must understand which compliance standards you need. Health tech companies will likely need to comply with HIPAA, including two-factor authentication, and data encryption. Companies in North America may want to get SOC 2 Type 2 accreditation to seek enterprise companies as clients since larger organizations require this to ensure functional security processes.

Ensure your chosen vendors meet regulatory requirements when selecting a cloud provider for metrics and other observability tools. Since observability providers will be storing your data by nature, they must follow your exact requirements. If your needs may change, choose a carrier with a considerable [breadth of regulatory accreditations](#) so your company can grow without needing to spend the extra time, money, and engineering to change vendors.

ALERTING AND ALERT TYPES

Alerting plays a prominent role in observability by telling the appropriate person or team in your organization when there is a problem. [Alerts](#) will send notifications when your metrics and associated monitoring system detect a threshold for your system. Without alerting, teams must spend time checking metrics to ensure the system is in good health. Correctly set alerts give peace of mind that the system is functioning so developers can concentrate on other work like building new features.

As helpful as alerts can be, they are only useful if your team can utilize the information. Alerts should be designed for quality and not quantity; they should be actionable and sent in a manner that will not fatigue your team. There are several different alert types available from various observability providers, and each type may have benefits for your service.

What is Alert Fatigue, and why should it be avoided?

Before deciding on the most useful alert for your system, consider the people receiving these alerts. Every time an alert comes in, a person will receive it and take action. For that person to just find the problem's source will take time, so selecting the best solution that will be most effective while also saving the most time is key.

[Fastly estimated](#) that 45% of all alerts from API security tools were false positives in 2021. The time taken to track such false positives will add up quickly. The same survey showed that 91% of companies turned off alerting and ran in log or monitor mode to avoid seeing any alerts, slowing their response time to actual threats. Alerts are essential to reacting very quickly to issues like these security threats, but the tools need to work for you and not against you. [Dynamic alerting](#) can help reduce false-positive alerts. This reduction helps companies [reduce their resolution time for errors](#) and improves the efficacy of engineering teams.

Static versus Dynamic Alerting

Static and Dynamic refer to how the monitoring tools set the thresholds to alert teams of problems. A static threshold will use the same threshold value to send alerts. Your monitoring tool alerts when it detects a value outside of your chosen threshold. Static thresholds are helpful in some cases, like detecting a change in a metric or checking for status code responses, like 200 or 404 responses from an API.

Dynamic Thresholds are more useful for data-stream monitoring. Dynamic monitoring will analyze patterns in your metrics to determine normal values for your system. These normal values adjust alert thresholds, so you are only notified when something is out of the ordinary. Dynamic alerting is one way to avoid alert fatigue amongst your team.

Setting a static threshold is not straightforward since a threshold may change with cyclic variations like time of day or during scheduled maintenance. Static alert thresholds can have over 90% more false positive alerts than dynamic alerting for data stream monitoring. All of the observability tools compared include static alerting.

Some of the compared tools also include dynamic thresholds that can adjust to patterns in your data. This includes functions like *outlier* that will find when a metric value is unusual using moving averages and standard deviation.

The best observability providers' features include [dynamic monitoring](#) to tune your threshold to account for patterns automatically. You do not need to set up specific metric thresholds; they are automatically calculated using machine-learning algorithms developed to enhance the observability of your solution. You can simply start sending metrics to the observability platform and gain insights before adding any customizations.

Specialized Event Alerting

Observability offerings with built-in machine learning can alert on specialized event types. Some providers have [built-in alert types](#) specialized to find certain event types within your metrics and provide more insights than static alerting. They can also provide general-use specialized alerts that can be used to trigger alarms on a variety of metrics. These include time-relative alerts as well as [four other](#) threshold calculation types.

[Time-relative alerts](#) are offered as a specific, unique alert type. These alerts are triggered when a ratio of events from one time to another reaches a pre-set threshold. These alert types enable DevOps teams to detect abnormal behaviors that other monitoring types may not detect. Time-relative alerting is useful for cases such as detecting when conversion rates increase (or decrease) relative to what they were on the same day the previous week or month.

DEFINING DYNAMIC METRICS

Software companies have long embraced logging to monitor their software and product solutions. According to the [VMWare State of Observability report for 2021](#), most companies use traditional or specialized logging techniques. The same report also noted that 82-92% of stakeholders agreed that cloud applications would have better availability performance if teams had visibility of application metrics.

Further to the performance enhancements metrics add, log data can be bulky and hard to view at a glance. Observability tools have created ways to convert logs into metrics automatically to optimize storage space and time to analyze data.

Types of Metrics Available

Metrics from logs can be thought of in two ways. First, you may want to know the count of logs arriving with a given format or value. For example, how many logs arrived indicating a specific service was initialized over time? Second, you may want to use the content of the log to calculate a metric. For example, if a latency value is present in a log, they are able to create an average over time and create a metric from this data.

Observability tools vary in their ability to derive metrics from logs; the [best tools](#) use Lucene in their tooling. The Lucene scripting language

gives more flexibility to what can be extracted and how the metric can be created. This is especially useful to set filters on other log values such as numeric status or subsystems.

VISUALIZATION OPTIONS

Visualization is a key to observability. The goal when visualizing your data is to get the answers and insights you need at a glance rather than needing to hunt for information. A good visualization sets observability apart from monitoring because it can allow your teams to react much faster to issues they would not see with monitoring solutions alone.

Most of the compared tools have some visualization tool either built-in or linked to the observability data. The only tool without built-in visualization is Open Telemetry, which is not an observability service, but rather a library that can help you standardize and collect observability data across your tech stack, avoiding vendor lock-in. All the other listed observability solutions do provide [visualization tools](#), but picking a tool that offers either their own UI or integrates with other tools is crucial. That being said, it's important to know, with both options, there are pre-built dashboards as well as the ability to fully customize them.

Proprietary Solutions include Custom Dashboards

A customizable dashboard is a must-have for any solution so that you can make the tool work for your system. Easy-to-use solutions provide [default data](#) out of the box on their dashboards so you have a starting point for your visualization. Dashboards then allow for adding customized graphs and widgets to your display. Widgets tend to be organized by visualization type first, allowing for displays to be line graphs, bar charts, heat maps, or lists. Vendors also differentiate widgets by how the data loaded will be defined. Widgets include query-based, pattern-based, parameter-based, and JSON values.

Plugins are also available for external visualization tools

Observability solutions will provide [plugins](#) that teams can use to migrate data to and from other tools. These plugins are especially useful if you already use these other tools, allowing you to add observability to your data to gain further insights.

Grafana is an open-source tool specialized in visualizing data. Some observability tools can be configured to send data to a self-hosted [Grafana display](#) to supplement or replace the hosted dashboard. The best tools can also display using a [hosted Grafana server](#), meaning you would not need to integrate your own Grafana instance for use.

A unique plugin also offered is [Tableau](#). Tableau is a commonly-used visual analytics platform. A common application of Tableau is for collecting marketing metrics. By allowing data to easily be sent to Tableau, you can combine your observability data for all metrics and get even more insights, such as understanding that a dip in conversion rate coincided with a platform outage.

PRICING

Costs tend to vary wildly across different observability vendors based on several factors including:

- Whether the software is proprietary or open-source
- Which level of service is required (basic, moderate, or enterprise)
- Which features are required (e.g., alerts, automatic metric creation)
- The amount of data your company will need to be reviewed

The listed cost of proprietary or the hidden cost of open-source?

Our comparison list of observability services for metrics includes four proprietary solutions and two open-source solutions. Prometheus is an open-source code base and has become a defacto standard for metrics. Open Telemetry is a more general tool maintained by the CNCF, which includes metrics along with traces and logs.

There is no cost to use Prometheus or Open Telemetry software since it is open-source. The cost associated with it comes from the engineering required to set up and maintain servers and the cost of server use. There are many options for using a hosted server from cloud providers like AWS, as it has managed solutions for both Prometheus and Open Telemetry. This removes much (but not all) of the engineering costs associated with using open-source tools but adds to the cost of the cloud provider. AWS charges for data ingestion, metric storage, and query processing.

Even if you choose to use Prometheus, other proprietary solutions could be used in conjunction using plugins. Tools will show how to [load Prometheus data](#) into their solution so the two could be used in concert.

Cost of Data Quantity and Internal Processing

Dynatrace, Sumo Logic, and Datadog break up their pricing model by feature set first, with 6, 8, and 14 feature set sections. Each feature set comes with a cost by GB, processing unit, or other company-defined units. An example of a company-defined unit is Dynatrace's *Digital Experience Monitoring Unit*, which calculates the cost of its Digital Experience Monitoring feature set. Feature sets include infrastructure monitoring, full-stack monitoring, cloud security monitoring, networking monitoring, and several more.

All of the proprietary solutions above are at least partially paid per volume of data. [Coralogix](#) exclusively charges for the level of analysis performed on the data. Users can use the TCO (Total Cost of Ownership) Optimizer to split the data into three pipelines (Frequent Search, Monitoring, and Compliance) based on their needs. They are the only proprietary solution not to limit features behind a paywall. Further, you can choose which pipeline your data is stored in to help reduce cost by storing infrequently-used information in a less-expensive category.

Since most companies break down their cost structure by feature set and volume, it is difficult to draw a meaningful comparison. To help estimate the volume you might expect for observability data, Instana ran an [experiment](#) for a simple Hello World application on a small Kubernetes cluster. The total volume of observability data produced was 452GB per month (with 285GB per month associated with metrics).

WRAP-UP

We have looked at why metrics are on the critical path to the observability of your software platform, what common mistakes are made in visualizing metrics, and who can provide tools to support your metrics. Metrics are one part of the path to observability and work best when also paired with logs, traces, and security for a complete picture of your software's health. When choosing an observability tool, consider the cost, regulations you must comply with, as well as that tool's power to provide you with the most useful information.

Monitor Your System Health at Scale

