

Jun 2024

INTRO TO OBSERVABILITY


By Manik Rana

Try Pitch



HI, I'M MANIK

- Background:
 - Mostly a web dev (React, Nextjs, Python, Go)
 - Student
 - Dogs are awesome



AGENDA

1. OBSERVABILITY

1. What
2. Components
3. Best Practices

2. MONITORING



1. What
2. Components
3. Best Practices

3. OBSERVABILITY + MONITORING

4. DEMO


BUGS ARE A RACE BETWEEN THE CUSTOMER AND YOU

- Who discovers the bug first?




BUGS ARE A RACE BETWEEN THE CUSTOMER AND YOU

- Who discovers the bug first?
 - If customer → your app is unreliable




BUGS ARE A RACE BETWEEN THE CUSTOMER AND YOU

- Who discovers the bug first?
 - If customer → your app is unreliable
 - If you → you provide great service!




BUGS ARE A RACE BETWEEN THE CUSTOMER AND YOU


- Who discovers the bug first?
 - If customer → your app is unreliable
 - If you → you provide great service!
- How do you stay ready to catch bugs?
 - Observability



SITUATION: AN ECOMMERCE APP




SITUATION: AN ECOMMERCE APP



DEGRADED PERFORMANCE

SITUATION: AN ECOMMERCE APP



What if we got
notifications for this?



What if we could track
all this?

What if we could know
whats causing it?

WHAT IS OBSERVABILITY

For a software application to have observability, you must be able to:

- Understand the inner workings of your application
-
- Understand any system state your application may have gotten itself into
-
- Understand the things above, solely by observing that with external tools
-
- Understand that state, no matter how extreme or unusual



WHAT IS OBSERVABILITY

A good litmus test:

- Can you understand what any particular user of your software may be experiencing?
- Can you determine anomalies even if you have never seen or debugged this particular state or failure before?
- Can you determine the things above even if this anomaly has never happened before?



To achieve observability, an application must generate specific data that helps in detecting issues and resolving problems.

and 3 types of data are needed:

LOGS

TRACES

METRICS

LOGS

Logs provide a chronological record of events or transactions within a system. They provide insights into when something happened and what exactly happened.

Tools like Elasticsearch, Logstash, or Loki are often used to deal with logs.


```
loki-1      | level=info ts=2024-06-21T17:56:25.164603251Z caller=table_manager.go:228 index-store=boltdb-shipper-2020-10-24
msg="syncing tables"
loki-1      | level=info ts=2024-06-21T17:56:25.165344334Z caller=table_manager.go:136 index-store=boltdb-shipper-2020-10-24
msg="uploading tables"
loki-1      | level=info ts=2024-06-21T17:56:25.165498209Z caller=index_set.go:86 msg="uploading table index_19895"
loki-1      | level=info ts=2024-06-21T17:56:25.165579501Z caller=index_set.go:107 msg="finished uploading table index_19895"
loki-1      | level=info ts=2024-06-21T17:56:25.165640668Z caller=index_set.go:185 msg="cleaning up unwanted indexes from tab
le index_19895"
loki-1      | ts=2024-06-21T17:56:25.165750251Z caller=spanlogger.go:86 level=info msg="building table cache"
loki-1      | ts=2024-06-21T17:56:25.166878459Z caller=spanlogger.go:86 level=info msg="table cache built" duration=1.005792m
s
loki-1      | level=info ts=2024-06-21T17:56:25.167838251Z caller=table_manager.go:271 index-store=boltdb-shipper-2020-10-24
msg="query readiness setup completed" duration=41.291µs distinct_users_len=0 distinct_users=
loki-1      | level=info ts=2024-06-21T17:56:25.170939751Z caller=table_manager.go:171 index-store=boltdb-shipper-2020-10-24
msg="handing over indexes to shipper"
loki-1      | level=info ts=2024-06-21T17:56:25.171051918Z caller=table.go:318 msg="handing over indexes to shipper index_198
95"
loki-1      | level=info ts=2024-06-21T17:56:25.171073918Z caller=table.go:334 msg="finished handing over table index_19895"
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:25.6922Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:25.8252Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:27.3292Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:27.3572Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:27.9542Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:29.6612Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:29.6922Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:30.4532Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:30.6032Z"}
app-1      | {"level":"info","message":"slow response","timestamp":"2024-06-21T17:56:31.0382Z"}
app-1      | {"level":"info","message":"This is an info log","timestamp":"2024-06-21T17:56:38.100Z"}
app-1      | {"level":"info","message":"fast response","timestamp":"2024-06-21T17:56:38.106Z"}
app-1      | {"level":"error","message":"This is an error log","timestamp":"2024-06-21T17:56:38.111Z"}
app-1      | {"level":"warn","message":"This is a warning log","timestamp":"2024-06-21T17:56:38.115Z"}
app-1      | {"level":"warn","message":"This is a warning log","timestamp":"2024-06-21T17:56:38.124Z"}
```

```
2020-04-12 12:52:43.313 +02:00 [Information] Engine initialized.
2020-04-12 12:52:43.428 +02:00 [Information] Supplied connection string will be modified to enab
2020-04-12 12:52:43.552 +02:00 [Information] Database already contains a table named "[dbo].[Sit
2020-04-12 12:52:43.615 +02:00 [Information] Supplied connection string will be modified to enab
2020-04-12 12:52:43.684 +02:00 [Information] Starting periodic task "ExpiredMessagesCleanup" wit
2020-04-12 12:52:43.686 +02:00 [Information] Starting periodic task "CleanupTrackedErrors" with
2020-04-12 12:52:43.705 +02:00 [Information] Database already contains a table named "[dbo].[err
2020-04-12 12:52:43.907 +02:00 [Information] Setting number of workers to 1
2020-04-12 12:52:43.915 +02:00 [Information] Bus "Rebus 1" started
2020-04-12 12:52:43.924 +02:00 [Information] Supplied connection string will be modified to enab
2020-04-12 12:52:43.940 +02:00 [Information] Database already contains a table named "[dbo].[Sit
2020-04-12 12:52:43.942 +02:00 [Information] Starting periodic task "ExpiredMessagesCleanup" wit
2020-04-12 12:52:43.943 +02:00 [Information] Starting periodic task "CleanupTrackedErrors" with
2020-04-12 12:52:43.957 +02:00 [Information] Database already contains a table named "[dbo].[err
2020-04-12 12:52:43.958 +02:00 [Information] Setting number of workers to 1
2020-04-12 12:52:43.959 +02:00 [Information] Bus "Rebus 2" started
2020-04-12 12:52:43.964 +02:00 [Information] Supplied connection string will be modified to enab
2020-04-12 12:52:43.978 +02:00 [Information] Database already contains a table named "[dbo].[Sit
2020-04-12 12:52:43.979 +02:00 [Information] Starting periodic task "ExpiredMessagesCleanup" wit
2020-04-12 12:52:43.980 +02:00 [Information] Starting periodic task "CleanupTrackedErrors" with
2020-04-12 12:52:44.000 +02:00 [Information] Database already contains a table named "[dbo].[err
2020-04-12 12:52:44.000 +02:00 [Information] Setting number of workers to 1
2020-04-12 12:52:44.001 +02:00 [Information] Bus "Rebus 3" started
2020-04-12 12:52:44.182 +02:00 [Information] SystemPerformanceCounters Constructor, Instance:Pro
2020-04-12 12:52:47.340 +02:00 [Information] Engine started
2020-04-12 12:52:47.377 +02:00 [Information] Sitecore Processing Engine has been started.
```

METRICS

Metrics are quantitative measurements that offer a snapshot of a system's performance over time. These include numerical values such as response time, error rate, or resource utilization.

Tools like Prometheus collect metrics from various sources such as nodes, Kubernetes services, and pods, and provide querying and alerting mechanisms.



TRACES

Tracing helps us to track the flow of requests through various services and components of a system

Tools like Tempo, Jaeger, and Zipkin help us track requests as they travel through various services in a microservices architecture.


Logs and metrics, when combined with traces, bring out better insights that help us to troubleshoot issues faster.

OBSERVABILITY BEST PRACTICES

- ✓ Decide what to log
- ✓ Cleanup logs
- ✓ The end goal of observability is that whenever there is an issue, we should have enough context to ensure that troubleshooting is effective enough to resolve any issue.

SITUATION: AN ECOMMERCE APP

Enter Monitoring



App



Search



Cart



Checkout



Payment

What if we got notifications for this?



What if we could track all this?

What if we could know what's causing it?


DEGRADED PERFORMANCE



Try Pitch
Very slow app. I don't recommend using this app!

SITUATION: AN ECOMMERCE APP

Enter Monitoring



- spot issues before they escalate
- monitor the average response time
- receive notifications when for a specified threshold.

MONITORING:
REDUCE DOWNTIME
INCREASE PERFORMANCE
PROVIDE BETTER USER EXPERIENCE

MONITORING:

REDUCE DOWNTIME

INCREASE PERFORMANCE

PROVIDE BETTER USER EXPERIENCE

Monitoring tools like Prometheus and Grafana help us to track the health of individual components including servers, databases, network devices, and other infrastructure.



MONITORING:

REDUCE DOWNTIME

INCREASE PERFORMANCE

PROVIDE BETTER USER EXPERIENCE

Monitoring tools like Prometheus and Grafana help us to track the health of individual components including servers, databases, network devices, and other infrastructure.

With monitoring, we can measure and consolidate key metrics of an application that are relevant to the business and actionable. These essential metrics can then be displayed on a dashboard and used to trigger alerts.



MONITORING:

WHAT KIND OF METRICS SHOULD WE MONITOR?

Monitoring everything is not a best practice, so to implement monitoring effectively, we need to know what to track.

MONITORING:

WHAT KIND OF METRICS SHOULD WE MONITOR?

Monitoring everything is not a best practice, so to implement monitoring effectively, we need to know what to track.

LATENCY


TRAFFIC

ERROR RATE

SATURATION


LATENCY

- Latency measures the time it takes for a request to travel from the client to the server and back.
- It measures how fast a response we are getting. Lower latency generally leads to a better user experience.
- In our case, when users search for a product, fast loading of search results is important. If search results take many seconds to appear, users start shifting to faster alternatives.




TRAFFIC

- Traffic denotes the number of requests a system receives over a specific period. Monitoring traffic ensures the system can manage load and scale appropriately as required.
- Again, in our case, during a flash sale event on our e-commerce web application, a significant increase in traffic can occur as users rush to make purchases.
- Monitoring this traffic helps ensure that the system can handle the load and scale appropriately whenever needed.




ERROR RATE

- The error rate metric represents the percentage of requests resulting in errors such as 404 page not found errors or 500 internal server errors.
- A high error rate indicates system instability or issues. Saturation measures resource utilization, including CPU, memory, and disk space.
- Heavy resource utilization can lead to performance degradation.



SATURATION

- Saturation measures resource utilization, including CPU, memory, and disk space. Heavy resource utilization can lead to performance degradation.
- For example, if a server's CPU usage remains near 90% due to extensive processing, it means it is saturated.
- Such high CPU usage might delay the processing of user requests, affecting the application's responsiveness.
- Monitoring saturation always helps us to prevent performance degradation due to resource outages.



By tracking these four golden metrics, teams can easily evaluate overall performance and application health, enabling informed decisions about optimization, scaling, and enhancements.

Without monitoring, we're operating blindly as there is no way to determine if the service is even functioning.

MONITORING BEST PRACTICES

- ✓ Monitoring should be adopted as soon as possible.
- ✓ We must take time to figure out what things we should monitor. Pay attention to the four golden signals.
- ✓ Dashboard and alert messages must be easy to understand and get right to the point.
If we send too many alerts, people might start ignoring them because there are just too many. So always limit the alerts based on priority.

So far, we learned that with monitoring, we can identify issues such as a specific service running slow or running low on CPU, and we get notified of such issues.

But getting these notifications alone isn't always enough.

Whenever any such issue arises, developers need sufficient data to fully detect the root cause and apply a permanent fix instead of applying some random fixes which waste time and resources. That is where observability comes into the picture.

OBSERVABILITY + MONITORING

- When an alert is received from monitoring, SREs or developers investigate using observability data.
- Both serve the goal of identifying problems and delivering a high-quality user experience.
- Monitoring and observability are different concepts that complement each other.

OBSERVABILITY + MONITORING

Key Differences:

- **Monitoring:**
 - Alerts us when something goes wrong.
 - Example: Alerts when the application response rate is slow.
- **Observability:**
 - Helps us understand why there is an error and how to fix it.
 - Example: Identifies which specific microservice is causing the problem, allowing for a permanent fix.

Practical Example:

- **Monitoring:**
 - Detects that application response rate is no longer fast.
- **Observability:**
 - Pinpoints the specific microservice causing the slowdown.
 - Provides the necessary data to apply a permanent fix.

ONE FINAL EXAMPLE

