



Reliable Software Systems

Week 3: Monitoring



January 24, 2018
Alyssa Pittman
University of Washington Allen School

Example outage: Instapaper

Had a 32-hour outage

One day, their database writes stopped working: "The table 'bookmarks' is full"

The database's underlying file system had a 2TB file size limit

The database had changed companies, been upgraded, etc

To restore service, they needed to explore solutions, dump and restore the database (took several hours)

<https://medium.com/making-instapaper/instapaper-outage-cause-recovery-3c32a7e9cc5f>

"If you can't measure it, you can't improve it."

-- Lord Kelvin? Peter Drucker? The internet?

”If you can’t measure it, you can’t detect or debug problems in it.”

-- me

What to monitor

Traffic

HTTP requests

API requests

Customer experience

Latency

Failures

Request status codes

Exceptions

Server utilization

CPU

Memory

Limits (e.g. open file handles)

Underlying causes

Application

CPU, memory

Garbage collection

Binary versions

And more...

What to monitor

USE method:

“For every resource, check utilization, saturation, and errors.”

Hardware: CPU, memory, network

OS: file handles

Software: thread pools, locks

<http://www.brendangregg.com/usemethod.html>

How to monitor

At development time

Design your monitoring and decide what to measure.

Instrumentation - add to your code & infrastructure to capture metrics.

At run time





System collects metrics as the application runs.

Centralized monitoring server polls your system to save them metrics.

Monitoring system generates dashboards and alerts.

How to monitor

Application servers

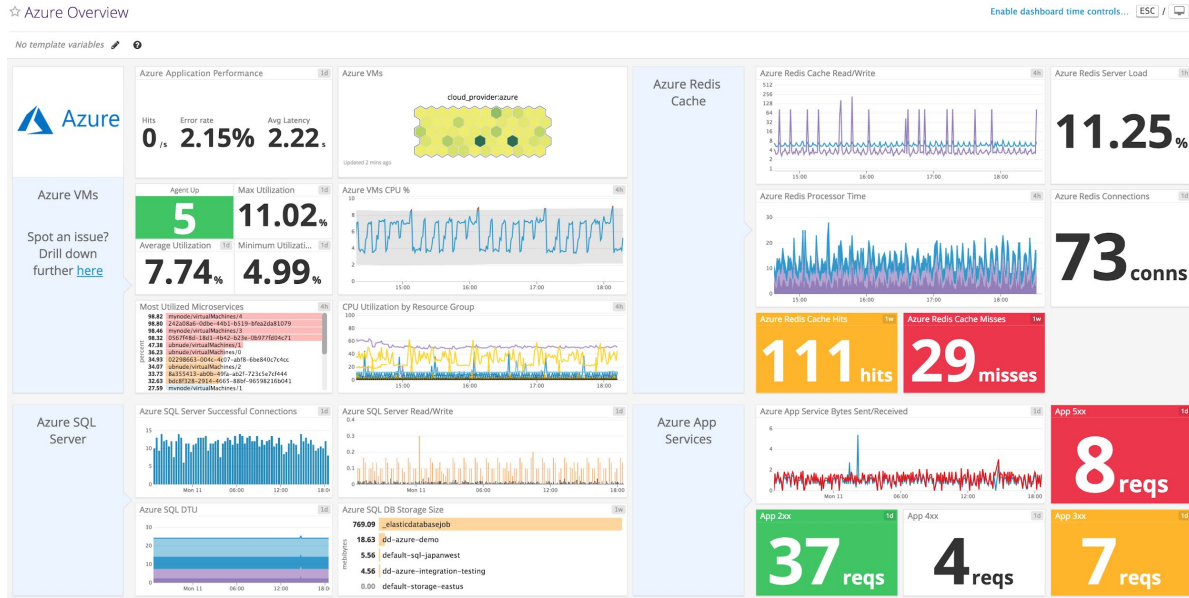


Timeseries database

Alerts

Data visualization

Demo (<https://www.honeycomb.io/play/>)



dashboard image from <https://www.datadoghq.com/blog/azure-monitoring-enhancements/>

Types of metrics

Counters - monotonically increasing numbers (e.g. time, total requests served)

Gauge - numeric value that can go up or down (e.g. temperature, memory usage)

Distributions (e.g. request latency)

Boolean (e.g. whether the cache was hit, whether a user was logged in)

String (e.g. an error type, a team name, a binary version)

They are recorded as timeseries - each value at a specific timestamp

Types of monitoring outputs

Dashboards

For visualizing, discovering, analyzing data

Alerts

For letting you know you have an urgent, important problem

Tickets

For letting you know you have potential, minor problems

Logs

For giving you detailed, textual information

Blackbox vs whitebox

Whitebox

You instrument the code to send metrics about inner details.

Examples: the metrics mentioned earlier, health checks

Blackbox

You use the system like a user and record the results.

Examples: [isitdownrightnow.com](https://www.isitdownrightnow.com), a prober

Metamonitoring



Sometimes necessary but use sparingly, only when you build/run your own monitoring, such as your own prober.

Information overload

Use frameworks and existing toolkits.

They often package some of these metrics for you.

Think carefully before introducing new metrics & logs.

Is this information already recorded somewhere else?

Will this be useful in an outage or will it overwhelm the rest of the data?

END

Alerting exercise

Service is receiving high QPS.

Service is receiving low QPS.

Server has high CPU load.

Service is responding with many failures.

A processing queue has a large number of messages in it.

A processing queue is taking too long to complete processing messages.

Database is almost filled to capacity.