



AN INFLUXDATA CASE STUDY


How Oracle Data Cloud Uses InfluxDB to Monitor Their Infrastructure

Tehmasp Chaudhri

Cloud Architect, Oracle Data Cloud



SEPTEMBER 2017



Company in brief

Oracle Corporation supplies software for enterprise information management. The Company offers databases and relational servers, application development and decision support tools, and enterprise business applications. Oracle's software runs on network computers, personal digital assistants, set-top devices, PCs, workstations, minicomputers, mainframes, and massively parallel computers.

A culture of innovation has propelled Oracle since its inception. Today, thanks to exceptional people, Oracle leads the market in autonomous, cloud, and applications technologies. Oracle's mission is to help people see data in new ways, discover insights, unlock endless possibilities.

Case overview

Oracle Data Cloud (ODC) provides the world's largest cloud-based data management platform for marketing teams. The Datalogix team of Oracle Data Cloud was tasked with providing managed infrastructure services to the Development team.

They chose InfluxDB platform because it served their metrics collection and storage needs and supported shifting their Dev team to containers. In addition, they built in monitoring to every container that their dev team spin up — making monitoring available from the get go.

Oracle Data Cloud uses InfluxDB to monitor their infrastructure; namely, the Mesos stacks that they give to Dev teams to do their workloads. Mesos stack monitoring using InfluxData enabled them to monitor cluster utilization and thereby control cost. In addition to InfluxDB for time series data storage, they also use Telegraf for metrics collection and Kapacitor for alerting.

"Beyond just collecting metrics, the big journey for us was not just giving Dev teams the ability to use containerization but really to give them a new platform for doing deployment of business services."

Tehmasp Chaudhri, cloud architect

The business problem

Oracle Datalogix, part of Oracle Data Cloud, needed to empower their dev teams by giving them a dynamic self-service delivery and development platform with integrated monitoring and alerting so that they could gauge performance and control cost.

The objectives they aimed at were:

- Provide managed infrastructure services to the company's Development teams
- Lessen the infrastructure burden on Dev teams to keep them focused on product features
- Migrate Dev services to containers to provide a future platform for development and delivery while finding Dev and Ops harmony
- Provide monitoring for infrastructure services (Mesos stacks) for free and automatically, by integrating their metrics and monitoring right into the stack — thereby empowering their developers to easily monitor the performance of their respective clusters as well as the applications they built within them

Upon deciding to provide their teams with Mesos clusters, they initially asked themselves the following business questions:

- If we do give Dev teams a cluster, will they be able to see a good cluster utilization for each of those stacks?
- When we provide these clusters out to the different Dev teams, how do we — as an ops team — see across the board, the total utilization for all those clusters?
- How are we capturing those cluster scaling events? Are we actually graphing that correctly and exposing that to the Dev teams?
- How do we reduce our cost by the visibility we can gain from seeing whether a given Mesos platform, on which we're running our services, has too many Mesos agents running?

The technical problem

The Dev teams they were helping either didn't have the ability to containerize their workloads or had been using containerization but taking on the burden of managing their own infrastructure. So the Datalogix team tried to provide those turnkey components for the business. Furthermore, they wanted to manage costs as well as help their developers monitor not just the performance of their clusters, but everything else the developer built within the cluster.

They needed a scalable, platform that would perform data collection at all levels of the stack to provide the flexibility and visibility that they and their Dev teams needed.

The solution

“InfluxData has been a great metrics tool and ecosystem for us at Oracle Datalogix and it really helps us to keep fostering and providing end-to-end infrastructure services for the organization.”

Why InfluxData?

Oracle Datalogix chose InfluxData’s platform to build their system of insight for metrics and events because it served their time series data storage, monitoring and alerting needs. InfluxDB – a high-performance data store written specifically for time series data – allows for high throughput ingest, compression and real-time querying. Not only can InfluxDB handle millions of data points per second, but it also offers flexible data retention policies. The reasons they cite for choosing InfluxData are:

- They found it to be, overall, a great metrics tooling ecosystem that has enabled them to provide end-to-end metrics for all their infrastructure services – including their containerization stack.
- InfluxDB has allowed the Datalogix Team to capture their own custom metrics to better answer some of their business questions around their journey in moving Dev teams over to containers.
- They found the team at InfluxData to be a great partner and the community to be very collaborative.
- InfluxData’s ecosystem continues to evolve and adapt to industry changes and viewpoints (for example, supporting push/pull models). They have been able to get feedback and help InfluxData grow its ecosystem which, in turn, helped the Oracle Datalogix team itself move forward as well.

The Oracle Datalogix team’s containerization stack

Oracle Datalogix provides the Dev teams with Marathon (a container orchestration platform for Mesos and DC/OS) and Singularity (an open-source project to manage computer clusters) on top of Mesos to allow these teams to containerize their workloads. The containerization stack they provided their Dev teams works as follows:

- They use a **templated Mesos cluster with heterogeneous slave pool support**. Each dev team gets their own cluster per cloud account. Teams can deploy multiple copies of the same application and run large data workloads in isolation depending on their needs. Before the Datalogix team began to support heterogeneous workloads, enabling their Dev teams to deploy different types of big data applications required extensive effort. Now, with heterogeneous workload support to Dev teams thanks to the monitoring integrated (to see how teams are utilizing the infrastructure), they can make the Mesos stacks as performant as possible.
- They deploy an **internal watchdog service** on each Mesos cluster to provide cluster elasticity and thereby control costs. The watchdog service monitors each individual Mesos cluster and looks for opportunities to scale the cluster up and down. They capture elastic event data and persist it back to InfluxDB so that they can visualize cluster utilization across time to their Dev teams.
- They capture as many cluster metrics as possible through **Telegraf plugins** and persist this data back to InfluxDB. They also capture additional cluster metrics via their own service that they deploy on each Mesos cluster and persist that back to InfluxDB.

Their containerization stack consists of:

- Mesos 1.3.x
- The latest Marathon
- The latest Singularity
- Marathon LB (a service discovery & load balancing tool for DC/OS)
- Zookeeper (which they try to manage as well as possible with Exhibitor)

Technical architecture

“We have two dozen Dev teams that we’re able to support at Oracle Datalogix. But initially, it was really hard until we added all those plugins with Telegraf to see exactly the visibility that was happening in each of those clusters.”

Below is a close look at the Oracle Datalogix metrics stack, covering each component of InfluxData’s open-source platform, in addition to how they use Grafana for dashboarding.

InfluxDB

The Oracle Datalogix team are running a single very large InfluxDB instance — 8 VPUs / 32 GB RAM — per cloud account:

- They have a 90-day default retention policy for base system metrics.
- They provide separate databases for Dev team application metrics (since they make the database instance multi-tenant), allowing teams to write their own data/retention policies into InfluxDB.
- They are working on migrating to InfluxDB Enterprise so that they can have a more spread-out workload and a more sharded setup.

Telegraf

For metrics collection, they have InfluxData's open source, plugin-based metrics collection agent, Telegraf, installed by default on every node of their stack.

- Initially, they used Collectd but switched to Telegraf since it had the plugins that they needed going forward to appease all the different infrastructure use cases at Oracle Datalogix.
- They have found Telegraf's plugin ecosystem to be extensive and performant.
- For their Mesos stacks, they install specific Telegraf plugins (such as Mesos, Docker, Zookeeper, and HTTP response) on their InfluxDB instance via config management discovery.

Kapacitor

They use Kapacitor as their main alerting engine:

- They templatize all of their TICKscripts so Dev teams can very quickly get some base alerting.
- Main alerts are set up and enabled on the same InfluxDB instance. Then they try to work with these teams to figure out what kind of other alerting paradigms they need.
- They provided their Dev teams with a container option for Kapacitor so that they can inspect their own metrics data in InfluxDB. They have also given Dev teams an ability to run their own Kapacitor containers to enable them to manage their own TICKscripts.
- They integrate Kapacitor with their chat and on-call system. This provides an end-to-end workflow — from metrics collection, to inspection into the database via Kapacitor, and then sending alerts that get triggered into the on-call system that they have set up.
- They view Kapacitor as a very good Swiss army tool of many uses, useful for both push and pull metrics.

Chronograf

Chronograf is mainly used by their infrastructure admins to monitor their InfluxData stack. They are waiting on more admin lockdown in the tool before assessing larger usage.

Grafana

They use Grafana as the main visualization system for InfluxDB metrics and for how their Dev teams interact with InfluxDB:

- Using Grafana, they pre-build template graphs for base system metrics and for Mesos cluster and utilization metrics.
- For Mesos metrics, they bring up the systems with their configuration management tools.
- Based on the cluster names that are represented to the Dev teams, the Dev teams are quickly able to visualize the metrics for their particular Mesos stack as well as the Docker container metrics for the specific containers running on their stack.
- This visualization setup has been very helpful in making dashboarding self-service.

They also pre-build template graphs for answering more interesting business questions such as:

- What is the cluster utilization of my Dev Mesos cluster?
- What is the total cluster utilization of all Mesos clusters provided by the Ops team?
- What are my cluster scaling events? Are we using too much money?
- Why isn't cluster scaling happening on this cluster?
- How many Oracle Datalogix services have we successfully helped Devs team migrate over to containers? Where's the real-time metrics/data?


Results

"The big onus on us was to provide these Mesos clusters out, but the nature of our workloads is so disparate that it's been a challenge for us to find the best generic use case. As we've been building out these clusters through all the metrics we've been gathering, we've been learning how to make these clusters more performant so they can serve many Dev teams."

The development and delivery platform that Oracle Datalogix built, and the integrated monitoring and alerting it provided, has helped them free their Dev teams from the infrastructure burden while providing visibility into Mesos cluster performance and utilization for both Dev and Ops teams. Their solution is helping answer the business question of whether they are making good headway into moving their Dev teams to containers and how well their Dev teams are utilizing Mesos clusters.

To demonstrate that newly gained visibility, below are some screenshots taken from their production side of using Grafana.

Mesos Cluster Graph



As shown above, Telegraf provides a lot of metrics. The graph visualizes one of their Mesos clusters and provides data from the Mesos admin GUI. It shows the number of killed and lost tasks, the error rates, and how many tasks in total are running on this cluster.

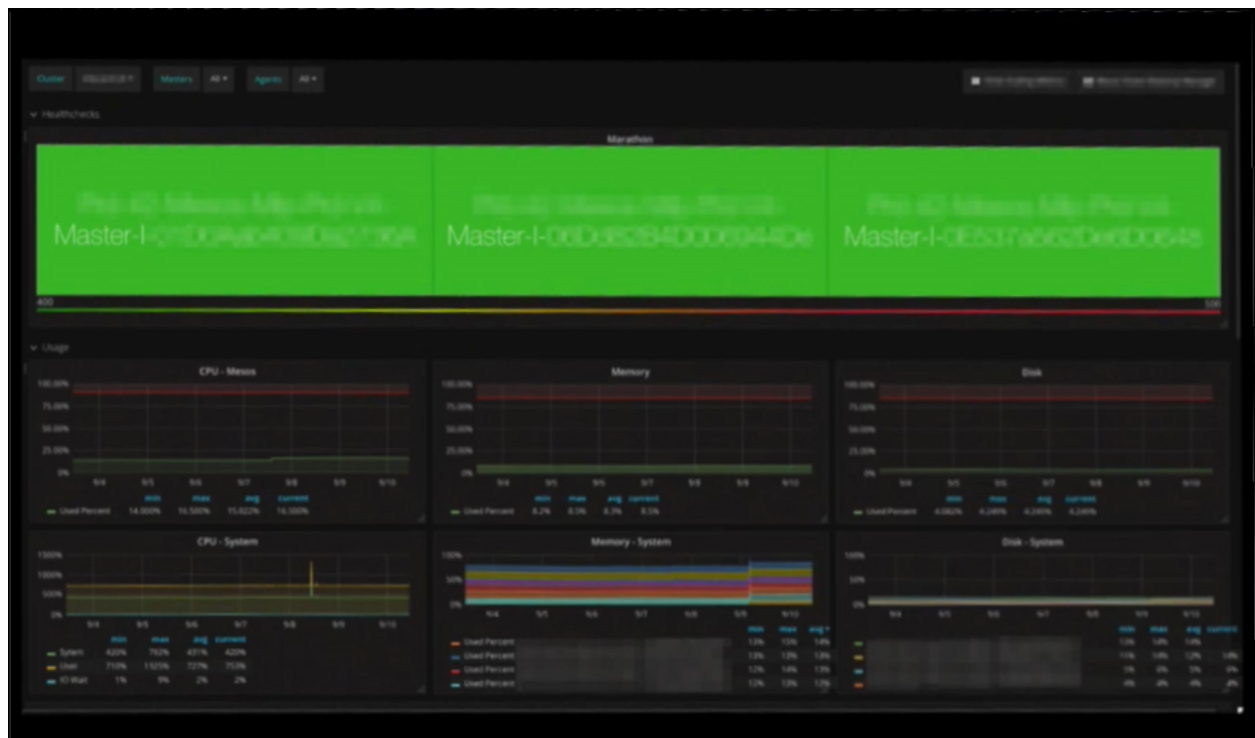
The below screenshot shows their Docker stats – all the different containers running on this particular cluster. This makes it very easy for any Dev team to see which Oracle Datalogix services are running on their cluster. The number of metrics that they collect from Docker using the Docker Telegraf plugin has been very useful as well.

Docker Stats Running on a Particular Cluster



Below is a dashboard that they built for showing the Marathon service healthtracks.

Marathon Dashboard



- The three green boxes in the dashboard above show that Marathon is operating very well.
- Below the green boxes, CPU metrics are graphed.
- At the top left corner is a cluster dropdown, enabling every Dev team to easily peek into their cluster and see the state of their containerization stack.

Additionally, outside of Telegraf, they have their own services running and feeding back data into InfluxDB. In this case, this is their own internal tool that they've been using to track all the Marathon applications running at Oracle Data Cloud. The below graph shows how they're building out a cleaner business intelligence dashboard which they can provide to the leadership team to see how many business services are running in containers.

Business Intelligence Dashboard



The chart below shows an incarnation of Oracle Datalogix's Mesos clusters in production. As shown in the chart's top left, in the totality of all the clusters that they provided out most recently, they were only utilizing those clusters globally at approximately 25%. So they plan to reach a higher cluster utilization across the board.

Oracle Datalogix's Mesos Clusters in Production



In their Dev workload, since they have more ability to do quicker jobs, they currently do see a higher cluster utilization. They have some Dev teams, due to the nature of their workload, being able to reach almost 60% utilization.

Cluster Utilization Stats



Using InfluxData, Oracle Data Cloud was able to free its Dev teams to focus on product development; meet its own needs for cost optimization; and ultimately to better achieve dev and ops harmony.

About InfluxData

InfluxData is the creator of InfluxDB, the leading time series platform. We empower developers and organizations, such as Cisco, IBM, Lego, Siemens, and Tesla, to build transformative IoT, analytics and monitoring applications. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by sensors, applications and computer infrastructure. Easy to start and scale, InfluxDB gives developers time to focus on the features and functionalities that give their apps a competitive edge. InfluxData is headquartered in San Francisco, with a workforce distributed throughout the U.S. and across Europe. For more information, visit influxdata.com and follow us [@InfluxDB](https://twitter.com/InfluxDB).



Try InfluxDB

Get InfluxDB

Contact us for a personalized demo influxdata.com/get-influxdb/

548 Market St. PMB 77953, San Francisco, CA 94104