Increase DevOps Throughput

How to Accelerate Release Frequency in Hybrid Cloud Environments

Summary

There is no doubt that speed is the new currency in our software-driven enterprises. For about a decade now, we have invested in automating the manual tasks in the software release process and maturing our CI/CD practices (or automating the pipeline). However, what you put through the “pipe” is just as impactful to speed as the automated pipeline itself. This is why managing data during the DevOps lifecycle has become a serious challenge to both DevOps throughput and agile velocity.

As organizations expand their footprint across the public cloud, while still managing on-prem environments, data management and provisioning becomes increasingly complex and potentially error-prone. This is why organizations need to take advantage of radical new advancements in data platforms. A data platform will transform the way your organization accesses data throughout the software development lifecycle—enabling greater DevOps throughput and optimizing cloud spend.

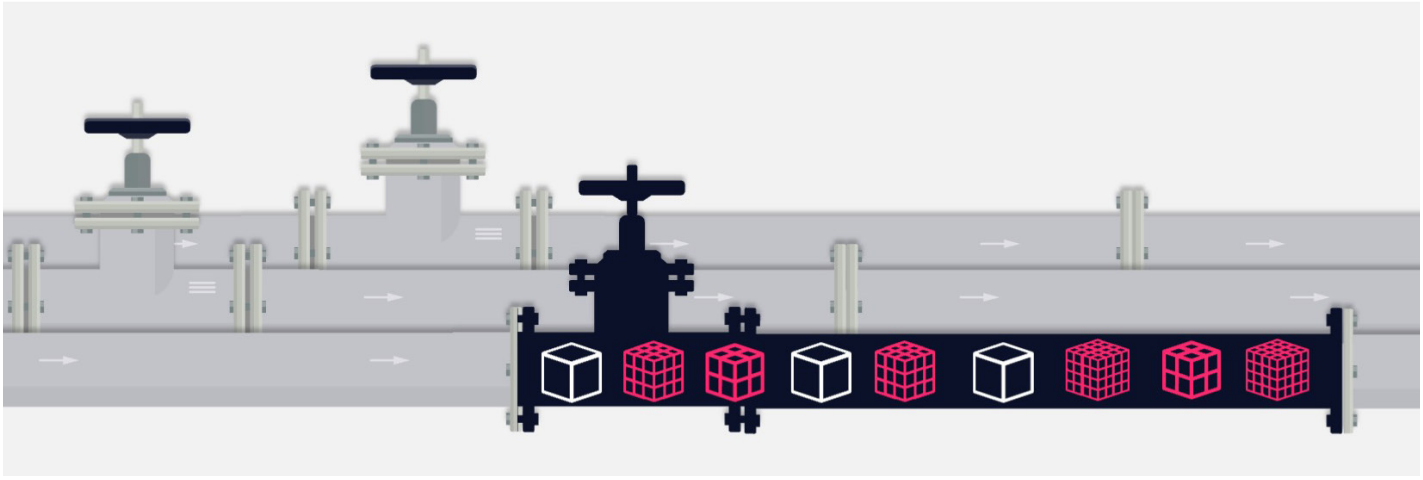

This paper explains how a data platform impacts DevOps by delivering real-time access to your complete data layer without refactoring for hybrid, multi-cloud environments. It also highlights how a data platform significantly impacts OPEX by eliminating manual tasks and increasing data efficiency.

The Challenge

Data management is a hot topic these days as organizations face a relentless increase in the type, volume, and platforms of data that need to be securely managed. This is all happening against the backdrop of DevOps, which is changing the way organizations build and use applications, both on-prem and in private and public clouds. The convergence of DevOps with explosive data growth has exposed a lack of sufficient automation of the data layer—creating the single greatest impediment to increasing DevOps throughput.

The prevailing method for managing changes in the data layer is Copy Data Management (CDM). CDM provides value when changing code for a single application within a single database, but it breaks down as an effective tool when being leveraged for a complex mix of legacy and modern applications that access multiple data sources simultaneously.

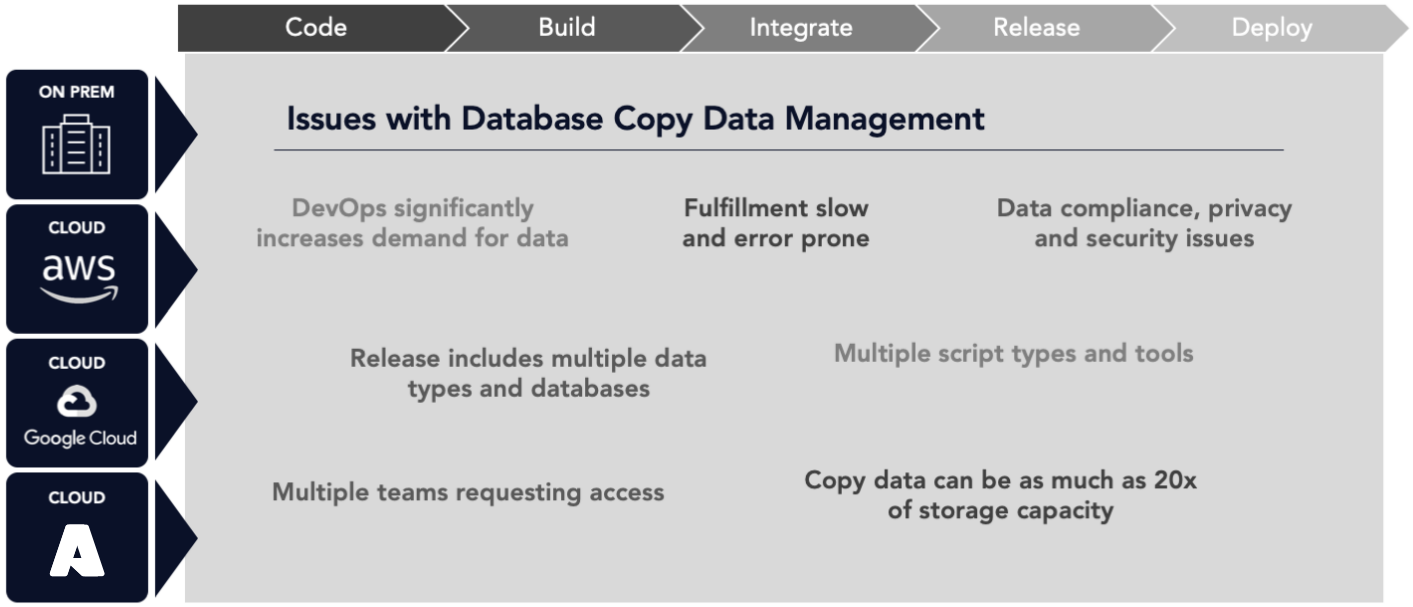

Using CDM to manage the data layer falls short for DevOps because:

1. The overhead of the CDM approach does not scale with modern architectures
2. DBAs are focused on stability, not change
3. Developers and testers make preliminary changes to the database independently of DBAs
4. Copy data grows logarithmically versus production data

The crux of the problem is that data management was not part of the original DevOps vision…and it needs to catch-up. The 2020 State of Database DevOps report found that 77% of developers work with databases and applications, while 75% of developers manage the database deployment scripts. The report also states that over 65% of organizations' DBAs are responsible for deploying changes to production. This highlights the friction in managing the data layer.

Perhaps the single greatest challenge facing a DevOps team is updating and optimizing legacy applications to a more modern architecture. Although there are best practices for a multi-cloud architecture, each public cloud is rich with nuances for optimizing an application within its domain. Cloud native architectures provide flexibility by introducing a componentized and distributed data layer. Although this componentization is optimal for DevOps throughput, it creates another challenge: it complicates the ability to migrate, maintain, and access data during Dev/Test across the hybrid environment.

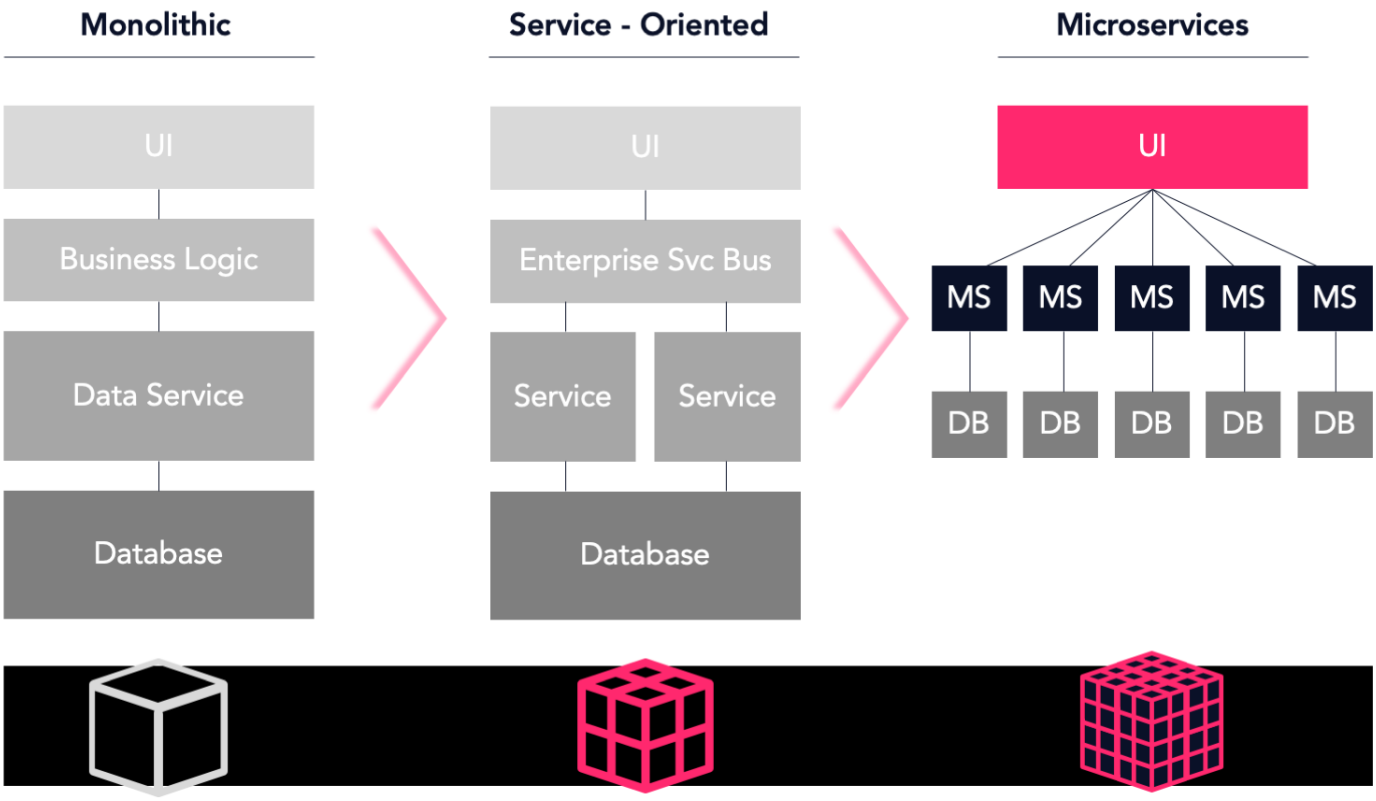

The advent of the public cloud launched a new breed of technologies focused on optimizing hybrid environments (on-prem and multi-cloud). By optimizing the data layer for hybrid cloud environments, a data platform abstracts and virtualizes your complete data infrastructure—agnostic of the brand or type of database in your enterprise architecture. This approach increases DevOps throughput and reduces cloud spend.

Get The Silk Platform technical details here.
Silk Platform – Mobilize your data for DevOps

The Solution – A Data Platform

A data platform fits between your full application stack and your data infrastructure. It decouples your application data from your on-prem hardware and public cloud infrastructure. This approach eliminates the risk of creating silos of data between multiple private and public clouds while providing a single access point for provisioning data across all stages of the software development lifecycle.

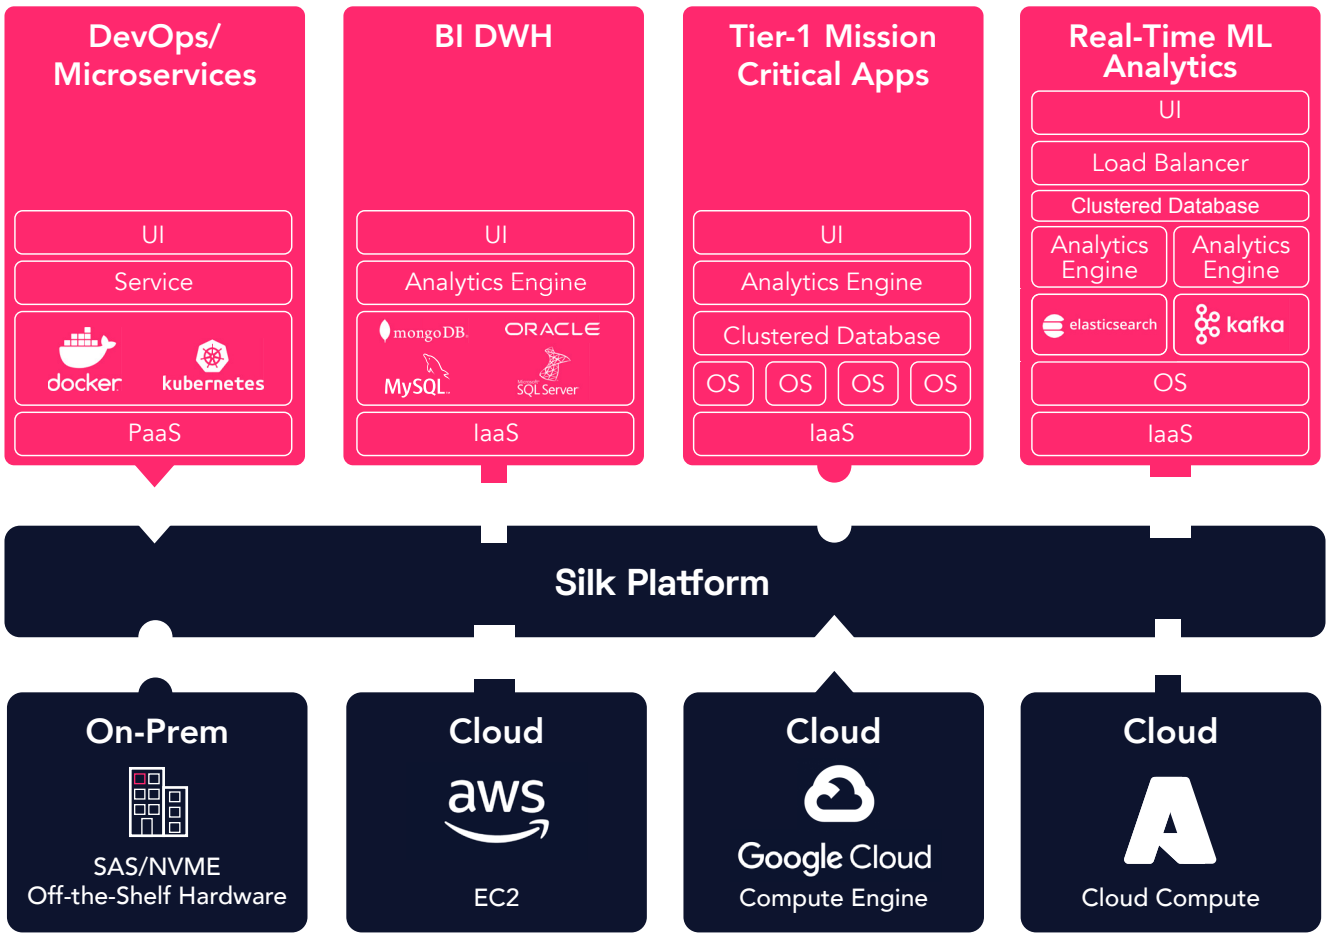

A data platform is a critical pillar in your DevOps strategy. It offers an unprecedented opportunity to eliminate wait times and avoid defects that stem from working with aged or error-prone copy data. A data platform is like adding a mesh wireless network to your home: it allows you to extend the value of your current infrastructure while giving you the flexibility to modernize how you work.

There are four primary drivers for leveraging a data platform to increase DevOps throughput:

1. Distributed agile development teams have a significantly increased demand for data
2. Testing modern applications requires access to broad and complex data
3. Refactoring for cloud is costly and time-consuming
4. Modern data platform optimize the data layer, resulting in a major reduction in cloud spend

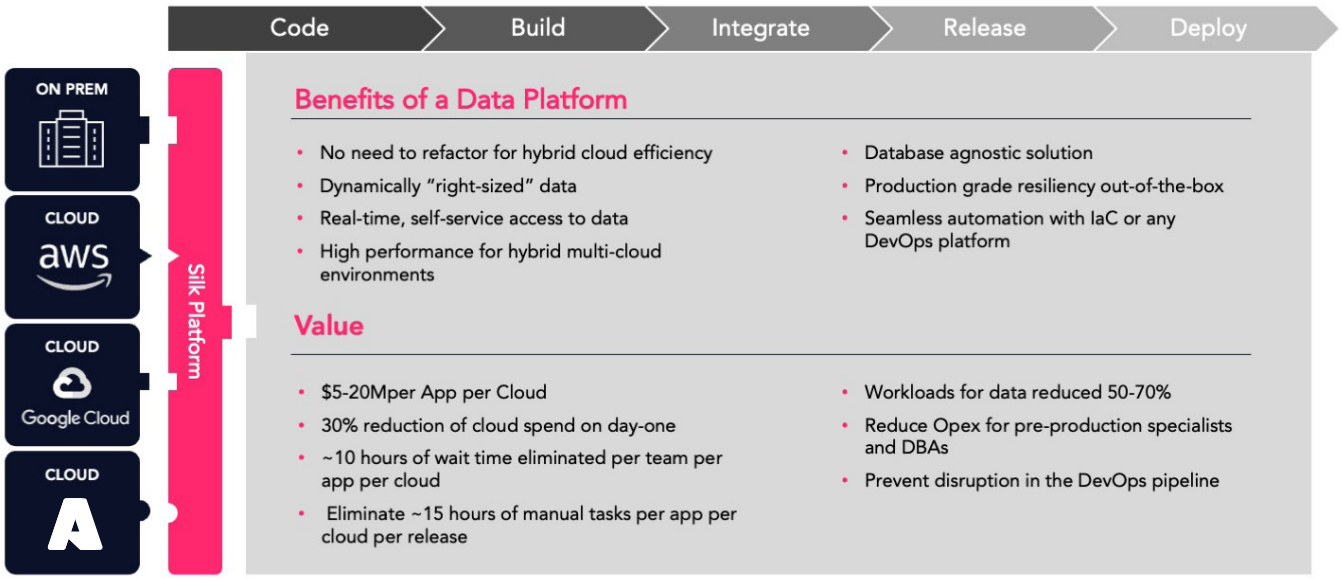

Distributed agile development teams have a significantly increased demand for data: As agile development methods expand in an organization, so does the demand for realistic, production-grade data. The team's velocity is dependent on data access, and the more distributed the team, the longer the wait time for data. A data platform eliminates wait times by providing real-time, self-service data access. Given the significant increase in the demand for both volume and types of data, a team can eliminate up to 10 hours per request.

Testing modern applications requires access to broad and complex data: With modern distributed architectures, functional and performance testing teams require access to significantly broader sets of data. The traditional manual request process brings release cycles to a full-stop. A data platform can eliminate this manual effort save roughly 15 hours per application per cloud instance per release.

Refactoring for cloud is costly and time-consuming: Refactoring the database is the most important and costly element when rearchitecting an application for the cloud. The data layer impacts storage, network, and compute. A data platformm abstracts the data layer from the application stack and optimizes for both speed and resource allocation. It allows legacy applications to run in a cloud environment without a "cloud cost penalty" and without the need to refactor. By abstracting the data layer, a data platform also optimizes DevOps throughput—speeding up release cycles.

Modern data platforms optimize the data layer, resulting in a major reduction in cloud spend: Storage accounts for roughly 35% of total cloud spend. A majority of this spend is attributed to running legacy applications (applications that have not been refactored or cloud-optimized) in the cloud. A data platform optimizes the data layer independently of the cloud infrastructure, taking the "legacy penalty" out of the equation. In turn, this results in a major reduction in cloud spend. See The Silk Platform for more detail.

About The Silk Platform

The Silk Platform helps organizations accelerate DevOps release cycles and reduce the risk of running business-critical applications in cloud-agnostic environments. With the boundaries between cloud environments blurring, Silk offers the ability to implement on-premise or any major public cloud platform, including AWS, Google Cloud Platform and Microsoft Azure. By using Silk, organizations can get ten times the performance out of their existing cloud data while spending 30% less. With real-time data reduction, thin-provisioning, and continuous resource optimization, Silk automatically matches cloud data spend to actual data needs.

Ready to get started? Visit www.silk.us to learn more about the Silk Platform.