# silk

Things to Know When Moving Oracle to the Cloud

Introduction

In an era where digital transformation is paramount, businesses like yours are increasingly transitioning to the public cloud to manage their data. The trend is not just about moving data; it's about leveraging the cloud's potential to enhance the most vital and sensitive parts of your IT infrastructure, particularly for handling essential workloads with systems like Oracle. The shift, however, is not without its complexities. Transferring extensive, intricate workloads such as Oracle to a cloud environment is a multifaceted process that presents a myriad of challenges. These challenges differ significantly from those encountered when migrating less critical workloads, which are often the pioneers in a company's cloud journey.

Our comprehensive eBook delves into the specifics of these obstacles and provides a roadmap for navigating through them to achieve a smooth migration of Oracle workloads to the cloud. Beyond just identifying the problems, we present actionable solutions and demonstrate the effectiveness of the Silk Cloud Data Platform. This innovative platform is designed to streamline the process, making the overcoming of these barriers migrating to the cloud more manageable. We go beyond theory by including a case study of a telecommunications company that has effectively navigated this transition. By sharing their story, we aim to equip you with the insights and confidence needed to embark on your own migration journey, ensuring that your critical operations are not just moved, but are transformed and optimized in the cloud environment.
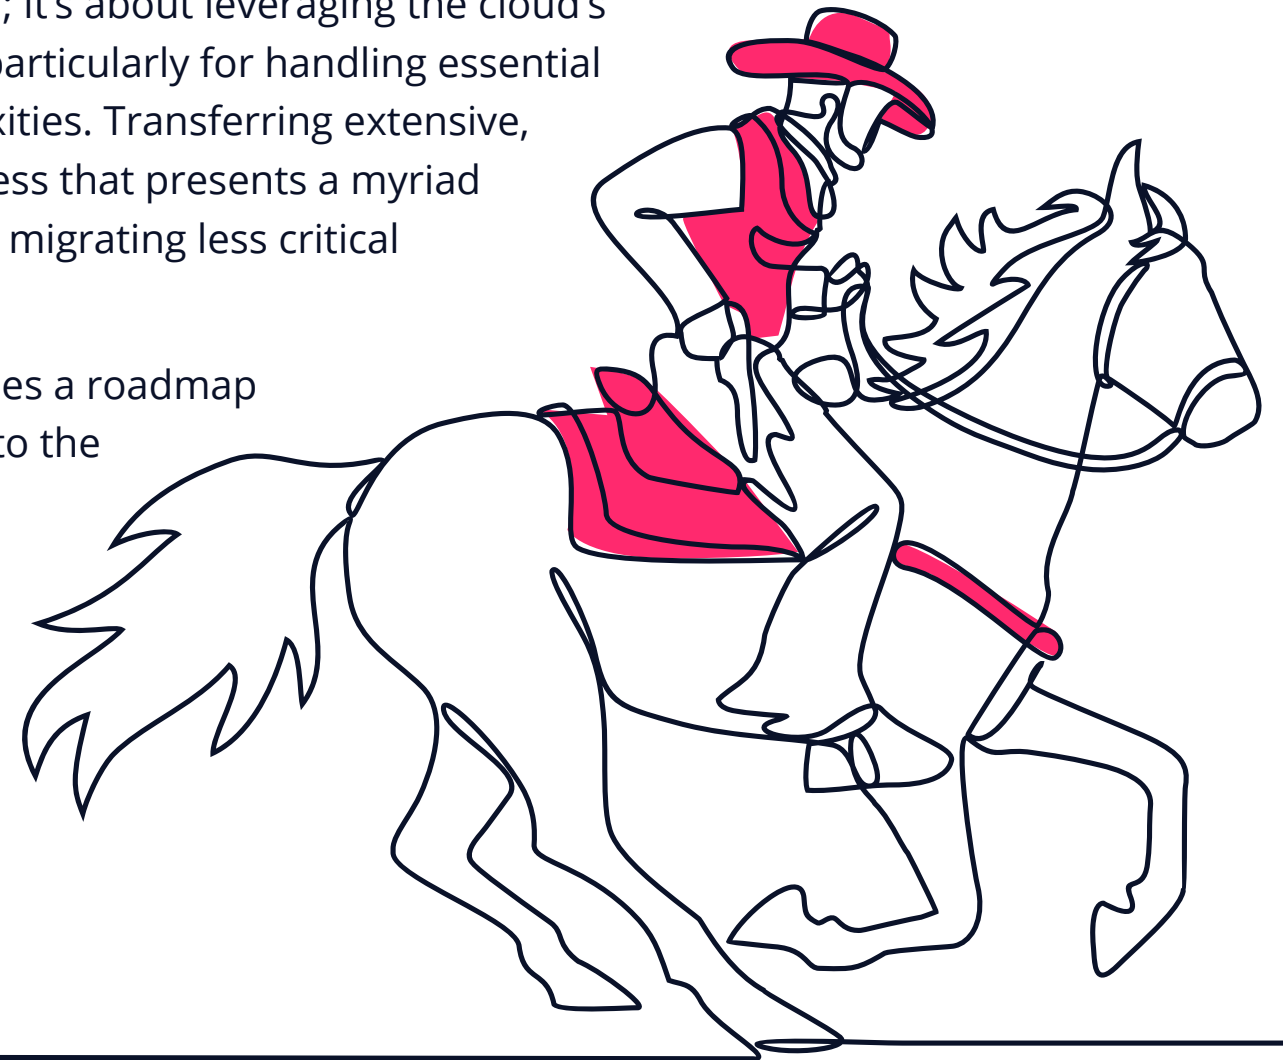
Obstacle 1: **Overprovisioning**

When organizations initiate the transition of their Oracle workloads to the cloud, they frequently encounter a critical planning obstacle: the necessity to overprovision their cloud resource allocation to meet demands. Overprovisioning is the process of reserving more cloud infrastructure resources — such as computational power, memory allocation, network bandwidth, storage capacity, and processing performance — than may be strictly necessary due to specific demands. This preemptive strategy is rooted in the interconnected nature of cloud resources, where enhancing one aspect often requires scaling up others, even if they are not immediately needed. This intrinsic linkage dictates a package upgrade approach, which can inflate the investment in cloud resources.

For Oracle workloads, which are particularly demanding at the storage tier, the requirement for ultra-fast performance is paramount. Achieving this in the cloud typically means committing to a considerable surplus of resources to ensure peak performance is available when needed. This can lead to provisioning far beyond your actual requirements and the surplus results in inflated operational costs as cloud providers will charge for the total resources provisioned, not just those in active use.

We will explore the intricacies of this issue, breaking down the mechanisms by which cloud billing can escalate due to resource overprovisioning, especially when migrating Oracle workloads. The following sections will dissect how Silk's data platform can mitigate these costs by enabling precise scaling and efficient resource utilization.

For those accustomed to on-premises Oracle systems, overprovisioning may be a familiar concept. Yet, the cloud environment accentuates this issue due to its inherent flexibility and the ease with which resources can be added on demand. This very flexibility, a significant advantage of cloud computing, can paradoxically lead to a lack of restraint in resource provisioning. Thus, without careful management and strategic planning, the cost-effectiveness of cloud computing can quickly become compromised, with cloud expenses spiraling unexpectedly. We will further discuss strategies to avoid such pitfalls, ensuring that your cloud investment is both efficient and effective.

Performance Data Analysis

Upon meticulous examination of the Automatic Workload Repository (AWR) report, we gain valuable insights into the Input/Output Operations Per Second (IOPS) and bandwidth that are necessitated by this system. The report delineates the demand for 210,000 IOPS for read operations and 37,000 IOPS for write operations. When we aggregate the data throughput, we observe it amounts to 210 megabytes per second for reads and 1.099 megabytes per second for writes. Such figures are crucial in understanding the raw data transfer capabilities required for efficient system performance.

IO Information from an AWR Report

| Function Name	Requests Per Sec- Reads	Throughput Per Sec- Reads	Requests Per Sec- Writes	Throughput Per Sec- Writes	vCPU Required	Memory for SGA/PGA
IO Stat Summary	52610	3459M	46234	989M	16 vCPU	204800M

In the above table, an AWR report for a significant high workload has been assessed to identify the needs for IO, vCPU and memory for the database to be migrated to the cloud:

vCPU: 16

Memory: 220G minimum

IO: 100K IOPS/4500MBPs

If we were to size out a VM using native Microsoft Azure IaaS solutions, the VM would need to be scaled, to the point that the Oracle licensing cost would be considerably more than the cloud infrastructure in the cost of the cloud migration project.
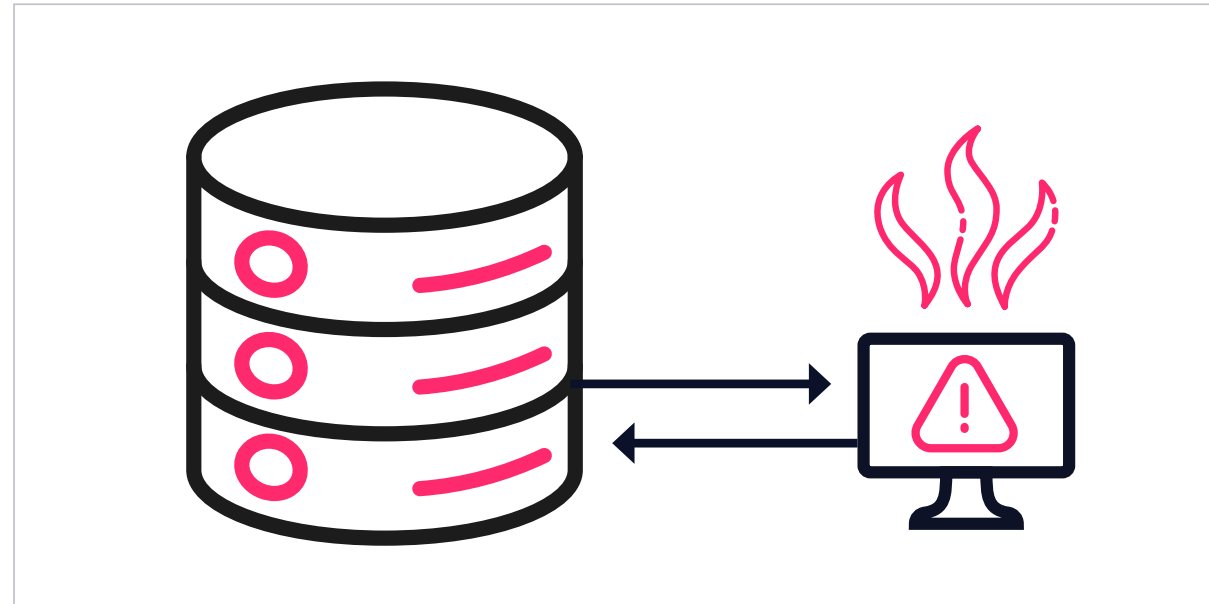
The core issue at hand is that most VM configurations provided by cloud vendors, even those supplemented with premium storage options, face challenges in meeting both the IOPSs and throughput demands of an Oracle workload without resorting to over-sizing CPU and memory resources. To achieve the requisite, I/O performance natively on Azure, the deployment of a VM with IO limits identified, along with memory and vCPUs, is imperative, as this often comes with significant cost implications., Pparticularly when considering Oracle licensing fees, which isare provisioned by each physical CPU. This consistently results in the majority of budget spent on Oracle licensing vs. the cloud migration project.This results in most budget spent on Oracle licensing versus the cloud migration project. Although the CPU might be adequately sized for the task, the financial impact of paying for unused cores can be substantial due to this fact.

Focusing on one of the few VMs able to handle this kind of high IO workload, the Ebs v5 series, we can quickly view the VMs that meet the needs of the workload:

Size	vCPU	Memory: GiB	Max data disks	Max uncached Premium SSD disk throughput: IOPS/MBps	Max burst un-cached Premium SSD disk through-put: IOPS/MBps	Max uncached Ultra Disk and Pre-mium SSD V2 disk throughput: IOPS/MBps	Max burst un-cached Ultra Disk and Premium SSD V2 disk through-put: IOPS/MBps	Max NICs	Network band-width
Standard_E16bs_v5	16	128	32	44000/1250	64000/2000	58960/1250	96000/2000	8	12500
Standard_E32bs_v5	32	256	32	88000/2500	120000/4000	117920/2500	160000/4000	8	16000
Standard_E48bs_v5	48	384	32	132000/4000	150000/5000	160000/4000	160000/4000	8	16000
Standard_E64bs_v5	64	512	32	176000/5000	200000/5000	160000/4000	160000/4000	8	20000

The table indicates that only a few VM options fulfill the requirements for vCPU, memory, and IO, necessitating an upgrade to at least Standard_E48bs_v5 to meet IO needs, which increases Oracle licensing costs significantly—up to threefold from the base price of $64,500 per processor license. With Silk, IO constraints are network-based, and features like compression and deduplication lead to better performance. This allows for use of the Standard_E16bs_v5 (16 vCPU, 256GB memory, and much higher limit with network IO), avoiding the need for a larger VM and extra Oracle licensing fees.
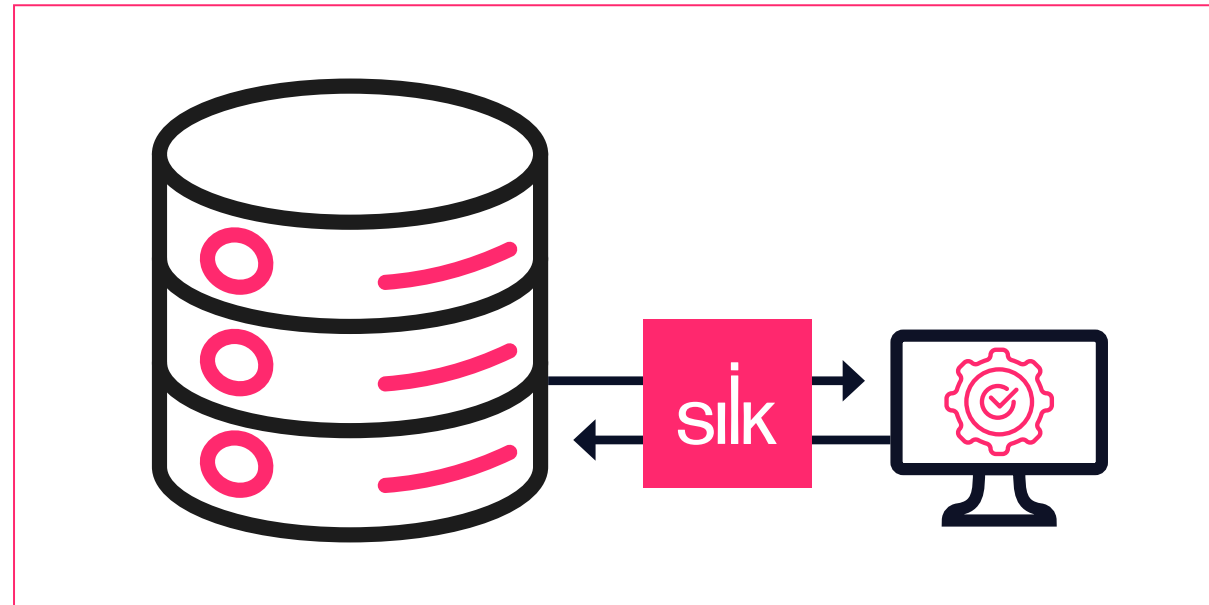
Silk's platform is instrumental in circumventing the dichotomy between vCPU/RAM capacity and I/O performance. It achieves this by leveraging a high-speed network bandwidth, which is significantly faster than the native Azure disk I/O pathways. By allowing for the independent sizing of vCPU/RAM and I/O dimensions, Silk facilitates a more precise and effective configuration of workloads. This targeted approach ensures that CPU resources are not left idle, and RAM is not overprovisioned, thus preventing unnecessary drains on IT budgets. By utilizing Silk, organizations can achieve a balance that aligns with their workload demands without succumbing to overprovisioning, thereby optimizing both performance and cost.

Obstacle 2
Increased Oracle Licensing Costs

When embarking on the journey of transitioning Oracle workloads to the cloud, one of the paramount hurdles that businesses must surmount pertains to the intricacies of core licensing costs. The labyrinthine structure of Oracle's licensing model requires meticulous attention, as it can ensnare the unwary in a web of unforeseen financial burdens. Oracle's Processor Core Factor Table is an essential tool for calculating the licensing requirements for Oracle's suite of products. Unfortunately it does not translate to the cloud environment. This omission introduces a significant layer of complexity, as the traditional benchmarks for estimating licensing needs are rendered obsolete in the cloud context.

The shift to the cloud further compounds this complexity due to the shared nature of cloud services. Cloud providers implement performance throttling measures to equitably distribute computational resources among their clientele. This shared economy model inherently imposes ceilings on the performance each customer can extract from their cloud services. Organizations may find themselves compelled to overprovision VMs with the intention of harnessing maximum performance. Yet, the reality is that even with this overprovisioning, the peak performance attainable is often a shadow of what is achievable on-premises. This diminished performance plateau comes with an additional sting—the necessity for a greater number of core licenses to support the beefed-up VMs.

A conservative estimate places the cost of Oracle licensing fees at 65% of their list price.

Oracle's licensing fees scale in a near-linear fashion, and they do so with considerable heft. A conservative estimate places the cost at 65% of the list price, and this figure does not even begin to account for additional features and capabilities that are often essential for enterprise operations. Options such as the Management and Diagnostics Pack, Advanced Security, Partitioning, or Advanced Compression all carry their own licensing fees. Given that these fees are tethered to the number of vCPUs, the cost can scale dramatically as additional computing resources are engaged.

It's vital to note that these additional features and options, while potentially increasing costs, are not mere luxuries; they are often critical components that enhance the security, manageability, and efficiency of Oracle databases. As such, their inclusion in the licensing calculus is not optional but necessary for maintaining robust and secure database operations in the cloud.

As a result, organizations must exercise a heightened degree of strategic planning and forecasting. The task involves not just a technical translation of on-premises workloads to cloud parameters but also a financial translation of licensing needs and costs. An in-depth understanding of both Oracle's cloud licensing policies and the performance characteristics of cloud service offerings is essential. Only then can businesses effectively align their Oracle cloud strategy with both operational requirements and financial constraints, ensuring that the move to the cloud does not inadvertently escalate into an unsustainable cost center.

Obstacle 3
Rearchitecting from On-Premises to Cloud Architecture

The transition to the cloud brings forth a multitude of challenges, one of which is the replication of on-premises Oracle features within the cloud infrastructure. For instance, enterprises utilizing Oracle's Real Application Clusters (RAC) on-premises benefit from the ability to distribute their database load across multiple servers. This horizontal scaling enhances performance and introduces a level of high availability critical for business operations.

However, this seamless horizontal scalability encounters limitations once enterprises look beyond Oracle Cloud Infrastructure (OCI) for their cloud migration. Oracle RAC, with its robustness and efficiency, finds itself in a unique predicament where its support is confined to OCI. Businesses looking to migrate to alternative public cloud vendors must therefore explore other avenues to replicate the availability, resilience, and efficiency provided by RAC.

One such alternative is Oracle Data Guard, which offers disaster recovery (DR) capabilities along with high availability. It also supports rolling patch updates, which are essential for maintaining uptime in mission-critical environments. Traditionally, RAC has been the go-to solution for horizontal scaling. However, the shift towards cloud environments has seen a preference for vertical scaling over horizontal due to its simplicity and the ease with which resources can be allocated. Oracle Active Data Guard emerges as a compelling replacement in this scenario, allowing organizations to offload read-only reporting and Recovery Manager (RMAN) backups to a standby server. This not only optimizes daily operations by utilizing the DR infrastructure for additional tasks but also ensures readiness for disaster recovery.
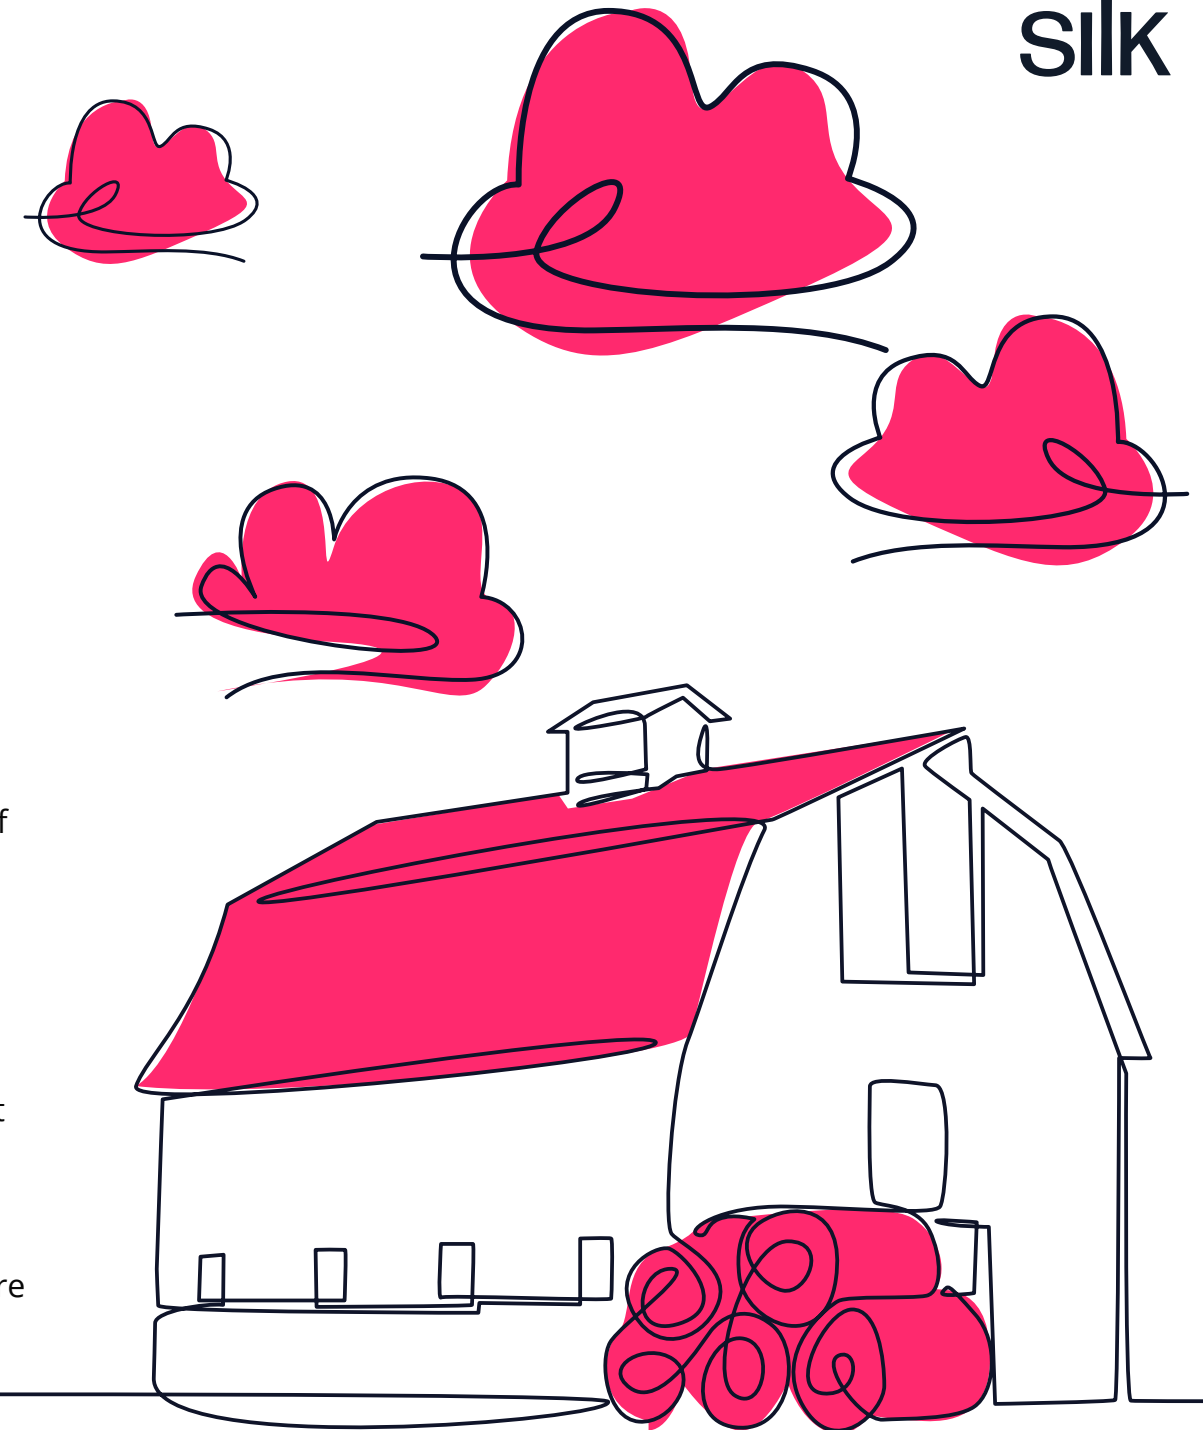
The goal of your cloud migration is to achieve a balanced, cost-effective cloud deployment that aligns the operational excellence of -on-premses Oracle systems, ensuring the cloud serves as a catalyst for business growth rather than a bottleneck.

Furthermore, Oracle's Hybrid Columnar Compression (HCC) technology, which significantly enhances data storage efficiency and performance by compressing and organizing data in a columnar format, is exclusively available for Oracle Exadata on-premises, ZFS storage, and Axiom Pillar SAN storage. The dilemma arises when seeking these benefits in the cloud since HCC is only available on OCI. To emulate the advantages of HCC elsewhere, cloud platforms must rely on alternative technologies that offer robust compression and deduplication capabilities to achieve similar levels of performance and storage savings. The `dbms_compression` package within Oracle systems can assist administrators in assessing the degree of compression applied and making the necessary adjustments for storage external to HCC.

In environments lacking HCC or equivalent compression and deduplication features, the repercussion is a noticeable expansion in data volume when migrating to other cloud providers. This data inflation inevitably leads to increased cloud storage costs. A workaround is the implementation of Oracle's Advanced Compression Option, which provides a means to mitigate the swelling of data. However, this option comes with its own set of trade-offs, including additional licensing fees. Furthermore, the additional computational power required often translates into the need for larger, more powerful VMs, adding to the overall cost.

The complexities of these scenarios underscore the need for a comprehensive approach to cloud migration—one that meticulously considers the replication of on-premises Oracle functionalities in the cloud environment. This approach should not only take into account the technical capabilities required to maintain performance and availability but also the strategic financial planning necessary to manage the total cost of ownership in the cloud. The goal is to achieve a balanced, cost-effective cloud deployment that aligns with the operational excellence of on-premises Oracle systems, ensuring that the cloud serves as a catalyst for business growth rather than a bottleneck.

How to Overcome These Obstacles with the Silk Cloud Platform

The journey of transitioning Oracle databases to the cloud environment is revolutionized with the integration of the Silk Cloud Data Platform. This innovative platform acts as a virtualization intermediary, seamlessly interfacing between the customer's Oracle workloads and the diverse array of underlying cloud infrastructures. It streamlines the process, facilitating the migration of even the most voluminous and intricate Oracle databases with relative ease. The Silk Cloud Data Platform effectively simplifies the 'lift and shift' strategy, which is often considered a daunting task due to the complexities associated with large-scale database migrations.

Silk's prowess lies in its ability to uncouple cloud performance levels from the volume of cloud resources utilized. It transcends the usual constraints imposed by cloud service providers, enabling optimal performance with a reduced resource footprint. By optimizing resource usage, Silk positions itself as a breakthrough solution that can offer, on average, a performance enhancement of up to tenfold compared to what would be achievable with native cloud configurations.

Moreover, the Silk Cloud Data Platform comes equipped with a suite of sophisticated enterprise data services. These services include advanced data replication techniques, the provision of instantaneous zero-footprint snapshots, and intelligent data deduplication methods. Together, they work in concert to refine the data management process, substantially reducing the overall data footprint and enhancing the efficiency of storage utilization.

A notable advantage of deploying Silk is the elimination of the need to overprovision resources. The platform empowers organizations to 'right-size' their cloud resource allocation, ensuring that they have just what they need to maintain optimal performance levels. This means that businesses can operate with leaner cloud resource configurations without compromising on speed or reliability, leading to cost efficiencies and more agile cloud operations.

What Does Silk Performance Look Like on Azure?

1M IOPS

20 GB/s
Throughput

Consistent
Sub-Millisecond Latency

One of the most significant benefits of Silk lies in its ability to minimize the financial impact of cloud migrations. By offloading a variety of database operations to the data layer itself, Silk mitigates the need for an excessive number of virtual CPUs on the database server. This, in turn, translates to lower Oracle licensing costs, as the number of required vCPUs—and consequently the licensing fees—are reduced.

The Silk Cloud Data Platform encapsulates all the advantages of Oracle in a public cloud setting. It is meticulously architected to eliminate single points of failure, thereby ensuring full system availability that withstands zone and region-wide outages. The platform's inline data reduction technology is particularly noteworthy, compressing and deduplicating data on the fly. This means that data storage is optimized without the additional burden of increased licensing costs. Moreover, the use of zero-footprint snapshots negates the necessity for creating voluminous full clones, which traditionally consume excessive storage capacity.

In essence, Silk provides a comprehensive solution that not only elevates the performance and resilience of Oracle workloads in the cloud but also does so in a cost-effective and resource-efficient manner. This convergence of performance, efficiency, and cost savings positions the Silk Cloud Platform as an invaluable tool for any enterprise seeking to harness the full potential of cloud computing for their Oracle databases.

Communication Provider Leverages Silk

A major residential telephone service provider in the United States embarked on a strategic mission to disentangle itself from the hefty licensing commitments of Oracle and Exadata. The objective was clear and pressing: to transition out of their existing data centers and migrate their technological infrastructure to Azure within a 12 to 18-month timeframe. The critical challenge they faced was the need to boost their workloads' performance beyond the capabilities of the native cloud, without undergoing the time-consuming and complex process of refactoring their applications—a process that was not feasible given the aggressive timeline.

In pursuit of a solution that could meet their performance needs while adhering to the tight schedule, the company turned to the Silk Cloud Data Platform. By leveraging Silk, they experienced a dramatic transformation in their reporting workflows on Azure—achieving a staggering 50% reduction in runtime compared to their previous Exadata on-premises setup. This leap in efficiency was not merely a marginal improvement but a transformative change that accelerated their business processes significantly.

Silk's suite of enterprise data services emerged as a game-changer for the company, enabling them to streamline their cloud resource utilization. With the enhanced data management and optimization capabilities that Silk provided, the company reduced the number of cloud resources required. This strategic resource optimization translated directly into substantial cost savings, as they were able to minimize their cloud expenditure without compromising on service quality or performance.

Moreover, the Silk Cloud Data Platform offered an unexpected benefit—the ease of cloud management that the company had grown accustomed to in their on-premises data centers. This familiarity was crucial as it minimized the learning curve and facilitated a smoother transition for their IT staff, who found that managing Oracle databases on the cloud with Silk was as straightforward as managing their legacy systems. This realization cemented Silk's position as not only a powerful performance enhancer but also a user-friendly and adaptable solution that mirrored the company's internal operational standards.
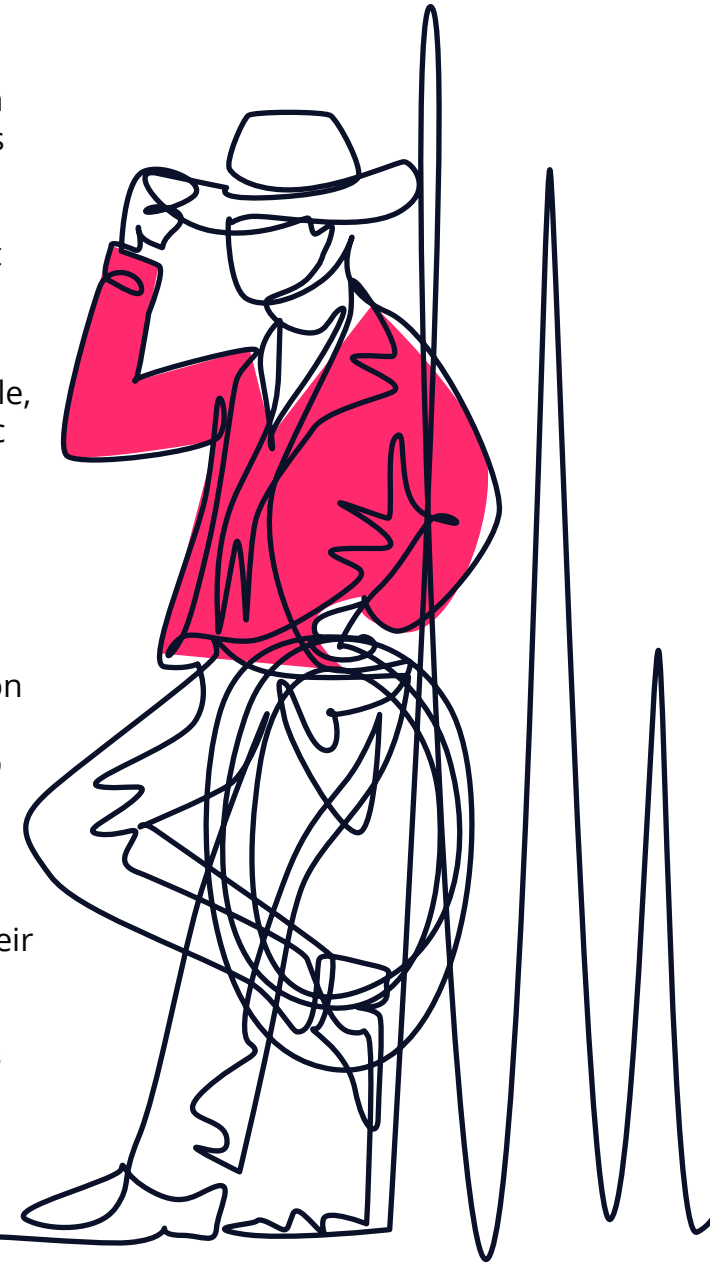
Silk demonstrated that it was not just a tool to lift and shift Oracle workloads to the cloud but a comprehensive ecosystem that enabled the company to replicate and even surpass their on-premises functionality. The company discovered that Silk was the scalable, straightforward solution they needed to not only meet but exceed their Oracle on the cloud requirements. With Silk, they were able to meet their strategic initiative well within the ambitious timeframe set, all while maintaining the integrity, performance, and manageability of their critical business systems.

Ready to move Oracle to Azure with Silk? | Learn more at www.silk.us