Moving Workloads From Oracle Exadata to Azure:

Top 5 Considerations

What's Behind the Trends

Most CIOs are still in the process of moving workloads from on-prem data centers into cloud environments, whether to close data centers, shift budgets from capital to operating expenses, or transform the company's ability to innovate.

But moving workloads from Oracle Exadata to the (non-Oracle) public cloud is no easy feat. Exadata is a clustered environment with multiple server instances working together to provide high availability, reliability, and scalability for Oracle workloads. Exadata is architected specifically for Oracle databases and mission-critical applications. And teams have often invested years building out the environment.

Migrating from Exadata to the cloud is challenging because it requires deep knowledge of Oracle, Exadata, and public cloud infrastructure. Exadata typically uses many Oracle features, such as Oracle Real Application Clusters (RAC) for availability and scalability and Oracle Hybrid Columnar Compression (HCC) for data compression, that are not supported on non-Oracle public clouds, so each must be evaluated and accounted for before making the move.

For Exadata on-prem customers, the road to cloud adoption is slow, costly, and full of risk – so it is no wonder that so many of these customers remain stuck on-prem today.

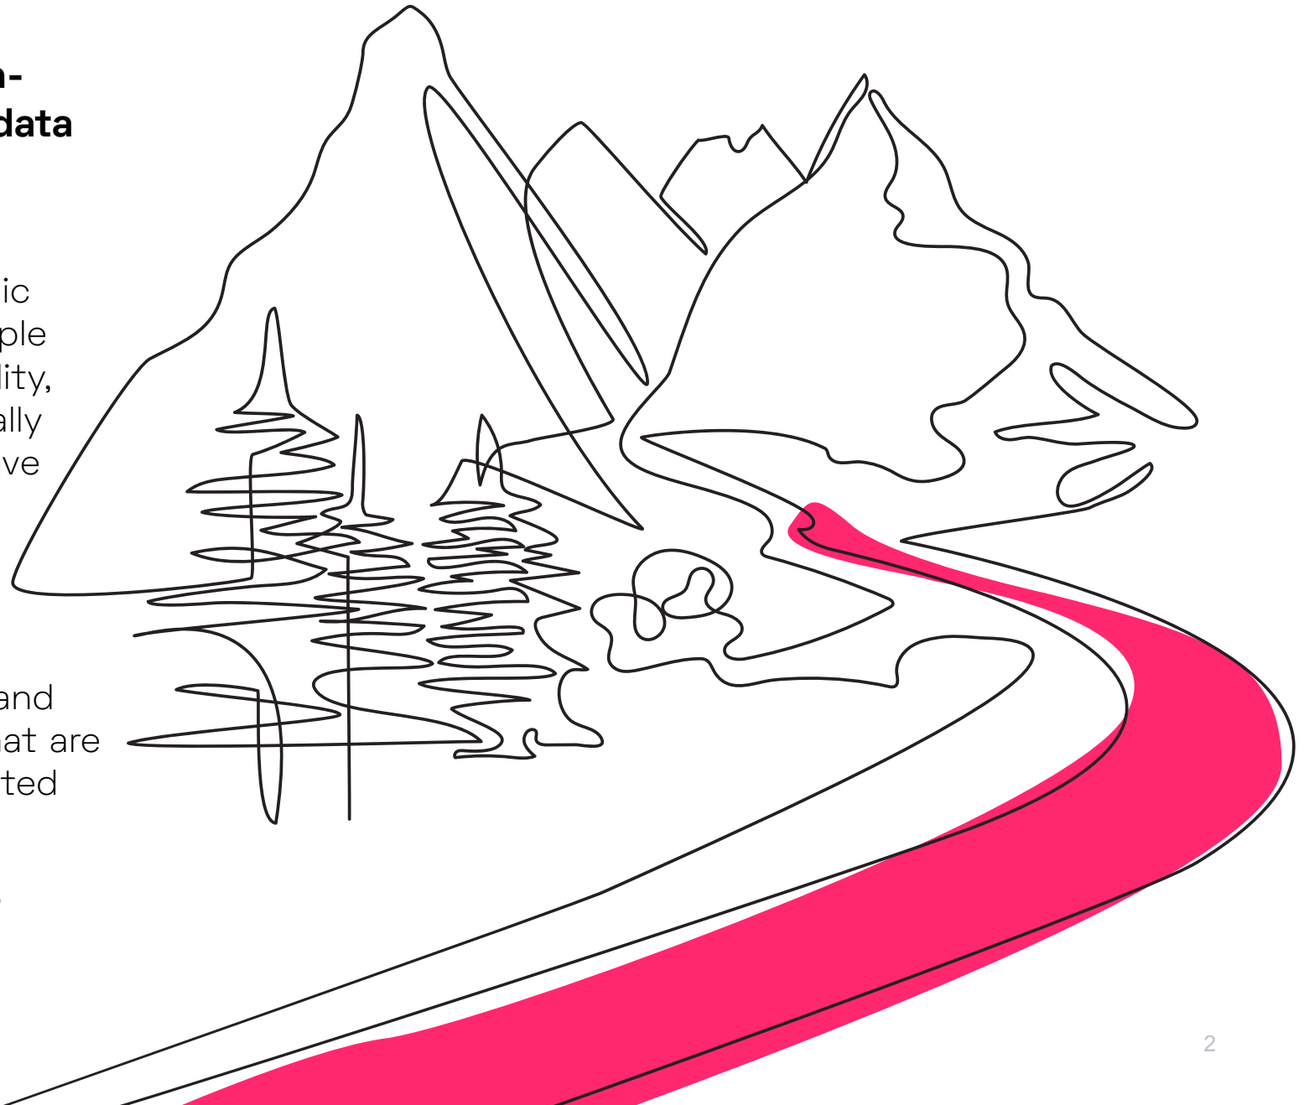

What's Holding CIOs Back from the Cloud?

This eBook looks at 5 considerations when moving workloads from Oracle Exadata to Azure:

1. Selecting the Right Migration Path
2. Achieving High Performance
3. Ensuring Data Resiliency
4. Enabling Cost-Effective Scalability
5. Cost of Cloud Services

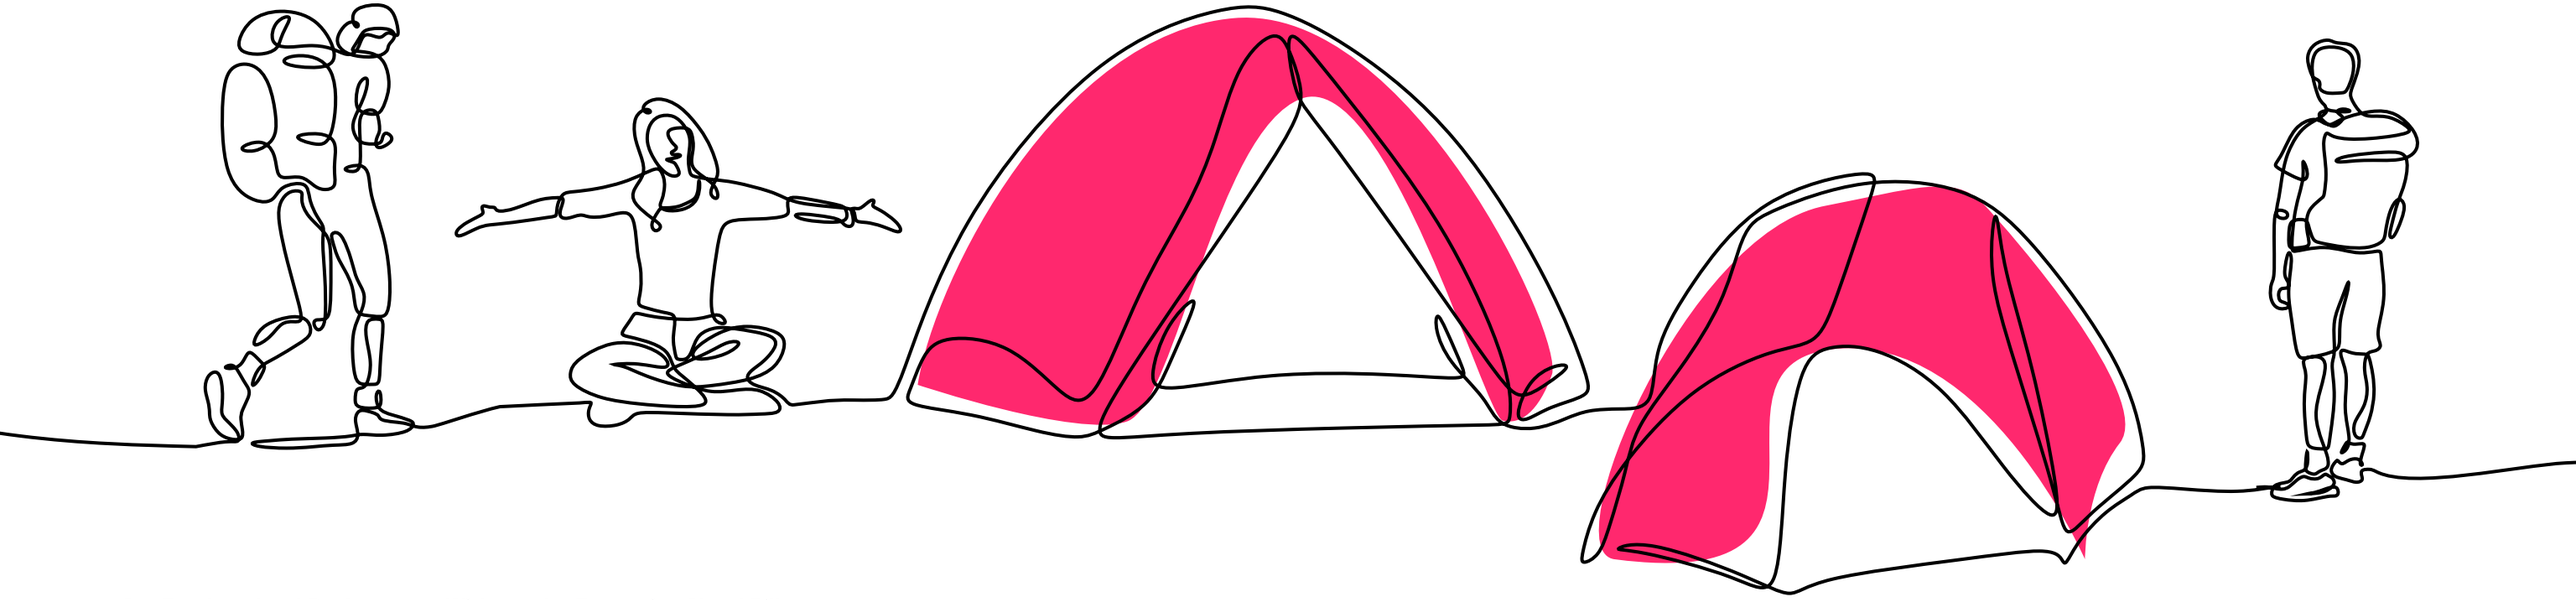

1. Selecting the Right Migration Path

There are two common options from cloud providers when moving Oracle workloads to the cloud – Infrastructure as a Service (IaaS) and Platform as a Service (PaaS). Each path has its own migration considerations.

With IaaS, an Oracle database and its applications run on virtual machines (VMs) on servers hosted on the cloud. No re-architecting is required; it's a lift-and-shift to the cloud. Setup and configuration are under the customer's control. This may be a good option for customers wanting visibility into and management of their databases, similar to on-prem. More importantly, application code does not need to be rewritten and business logic can remain untouched. But one drawback to IaaS is the inability to harness the infinite scale of the public cloud without a refactor. For example, since Oracle RAC is no longer supported, Oracle must run on a single-instance database host – which is not an ideal solution when needing to scale performance up and down.

With PaaS, also referred to as Database as a Service (DBaaS), the offering is a managed database service from one of the major cloud providers. Database management is taken care of, including updates, patching, security, and so on. This path often involves a database migration, typically from Oracle to PostgreSQL, depending upon cloud migration plans and application complexity. PaaS is more expensive than IaaS and often requires a refactor, unless the cloud provider offers an Oracle PaaS service. It is therefore less suitable for complex systems or where business logic is deeply embedded in database code.

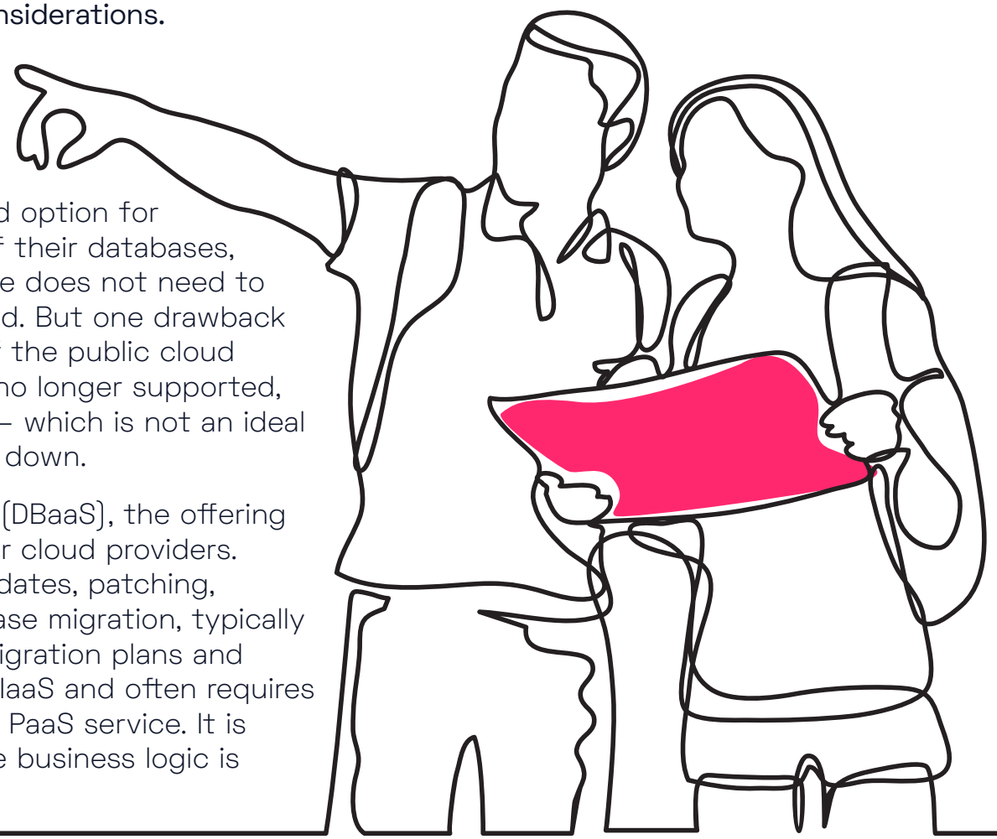

Additionally, each cloud provider offers different services and support for Oracle databases and applications.

- Oracle does not certify its software running on Google Cloud Platform
- Oracle software is certified to run on Microsoft Azure IaaS, but there is no PaaS option
- Amazon Web Services offers IaaS and PaaS solutions for Oracle, but the latter has limits around database size

Each migration path has its benefits and drawbacks. And every database and application have different needs. Therefore, evaluate each independently to determine the best path forward.

2. Achieving High Performance

Whether you're running online transactional processing or analytical workloads on Exadata on-prem, it's critical to achieve the same levels of performance on the cloud to meet service level agreements and customer experience demands.

In public cloud architecture, database performance and cloud capacity are inextricably linked. When a customer wants more performance, they need to take more cloud capacity (additional and larger VMs) and pay for it, regardless of whether it's used. Customers inevitably end up overprovisioning and oversizing VMs to hit mission-critical performance levels. The cost for performance can make it prohibitive to lift-and-shift workloads from on-prem.

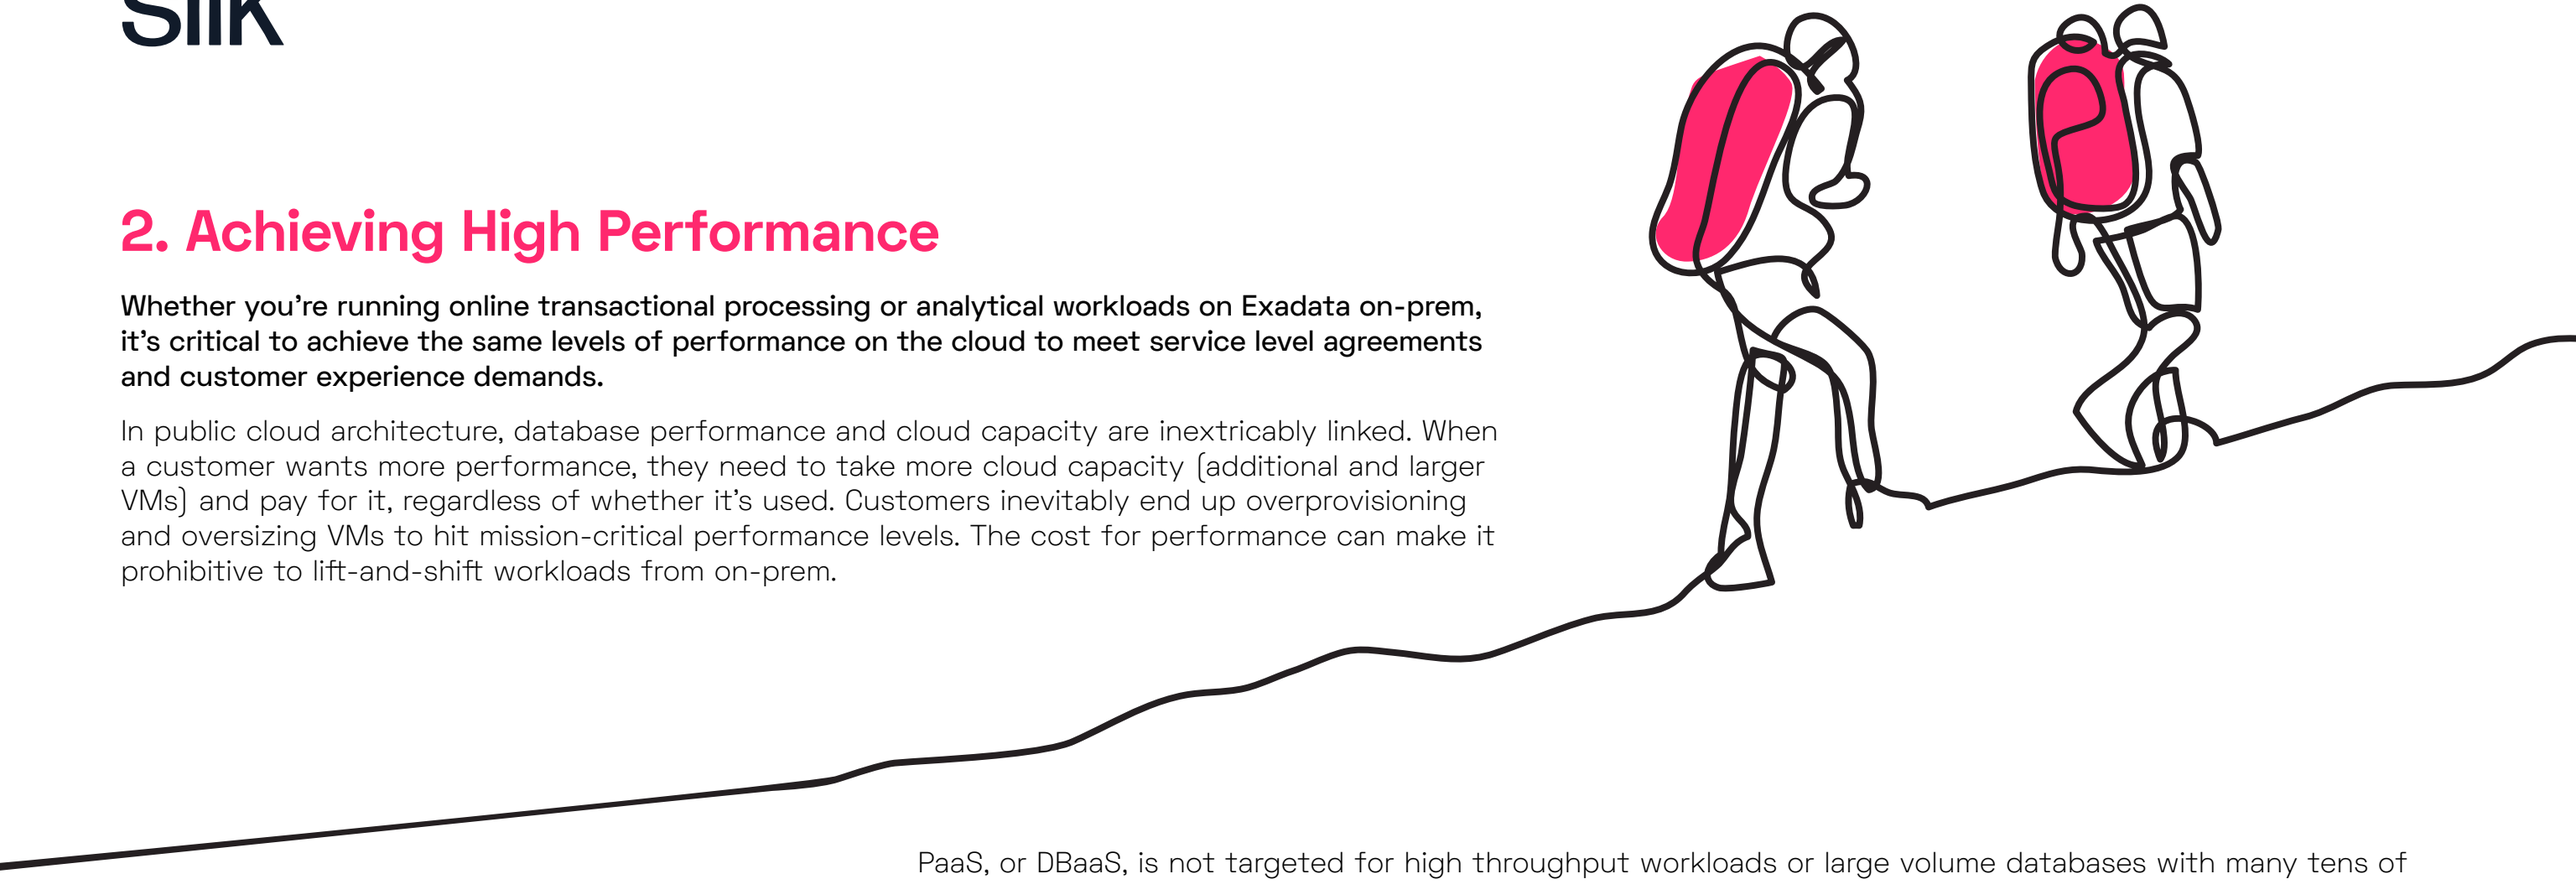

PaaS, or DBaaS, is not targeted for high throughput workloads or large volume databases with many tens of terabytes of data. When customers try to map such workloads onto PaaS, they can end up with oversized and overpriced PaaS offerings, with the need to distribute across multiple PaaS databases. Cloud bills get out of control.

For workloads coming to the cloud from Exadata, we highly recommend a cloud data platform, like Silk, that decouples database performance from cloud capacity to deliver significantly faster performance without needing to overprovision cloud resources.

3. Ensuring Data Resiliency

The database is the heart of many applications, and data resilience is critical to business operations. Moving workloads from Exadata to the cloud takes significant planning and the project may require re-architecting for resilience.

Many Oracle customers have used Oracle RAC to build highly available solutions across multiple nodes. Since Oracle RAC is not certified by Oracle on the non-Oracle public cloud, customers need to consider the impact on resiliency when moving from a clustered environment on Exadata to single-instance Oracle on the cloud.

One option is to use Oracle Data Guard instead. Through the use of synchronous log replication and transparent application failover, Data Guard can be configured to provide a zero RPO solution for single-instance Oracle. Data Guard is also included in Oracle's Enterprise Edition license, so it may be more cost-effective than RAC (which is an additional license option).

Additionally, the cloud providers offer various geographical regions and availability zones to help customers architect for resilience. Customers can choose the best regions and zones for their needs based on technical and regulatory considerations, such as latency and data residency.

Ultimately, data resiliency on the cloud is a shared responsibility. The cloud provider is responsible for the resiliency of their infrastructure, including the network and data centers. Customers are responsible for the architectural integrity of applications and workloads.

4. Enabling Cost-Effective Scalability

Performance demands will increase on the cloud, and customers want to cost-effectively scale cloud resources as this occurs. Yet there are certain limitations to cloud architecture that can hinder the economics of scalability. Performance of the data layer is tied to how much capacity and compute is provisioned, so customers can end up overprovisioning VMs by architecting to cover peak periods. There is waste in cloud capacity not being used during “normal” operations.

In another example, Oracle RAC is a technology used by customers for horizontal scalability. Oracle RAC enables the customer to add database nodes, which incrementally adds performance to the database cluster. When workloads are moved from a clustered environment on Exadata to the non-Oracle public cloud, where single-instance is the only supported option for Oracle, getting more performance equates to adding more compute capacity to the host virtual machine (vCPUs), which drives up the cloud infrastructure bill and Oracle license costs.

Customers must therefore consider how the cloud providers scale resources as demands for performance change.

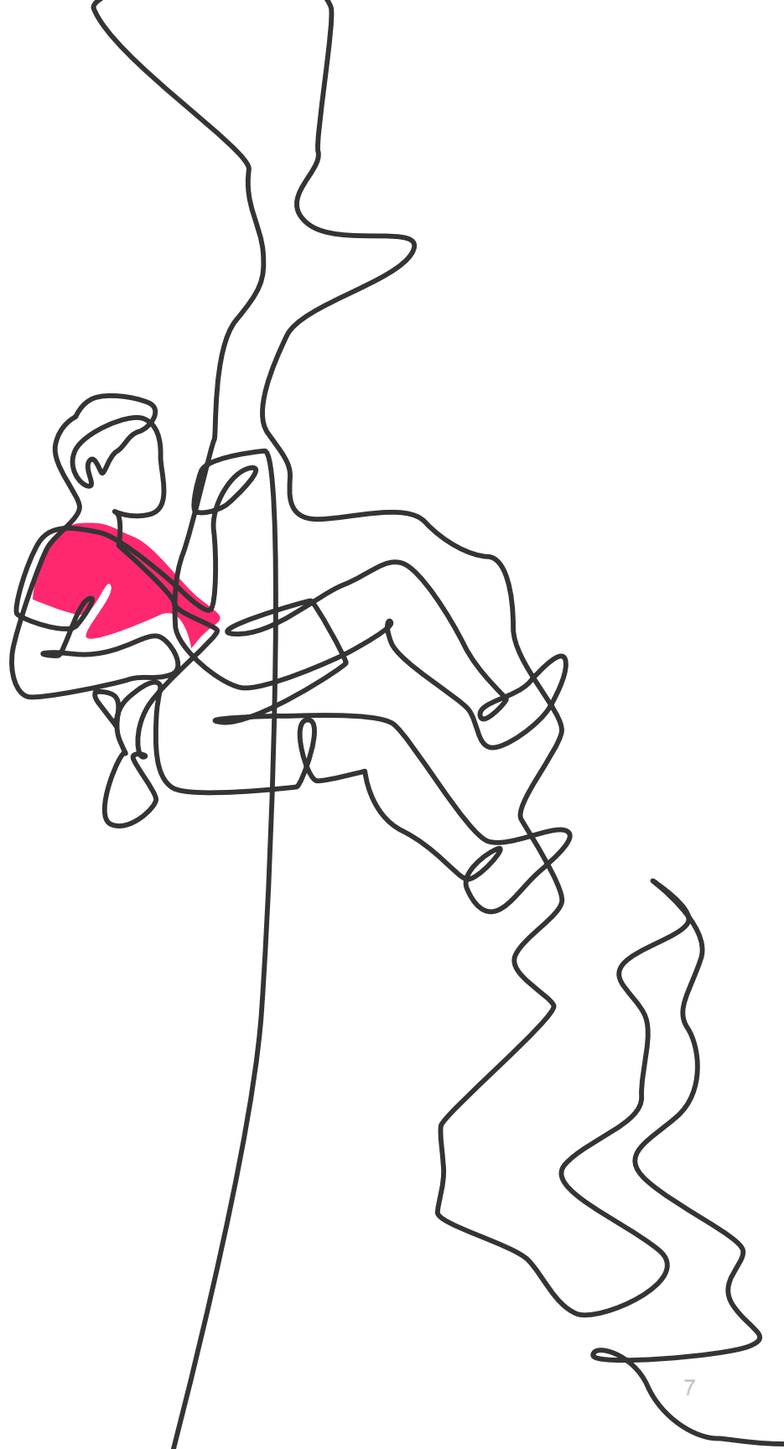

5. The Cost of Cloud Services

"How much is this going to cost me?" - every CIO on the planet. The team needs to determine the costs of preparing for migration, the migration itself (including comprehensive testing), and the running of workloads on the cloud. Armed with this information, the CIO can decide how to proceed and how quickly the cloud migration can begin.

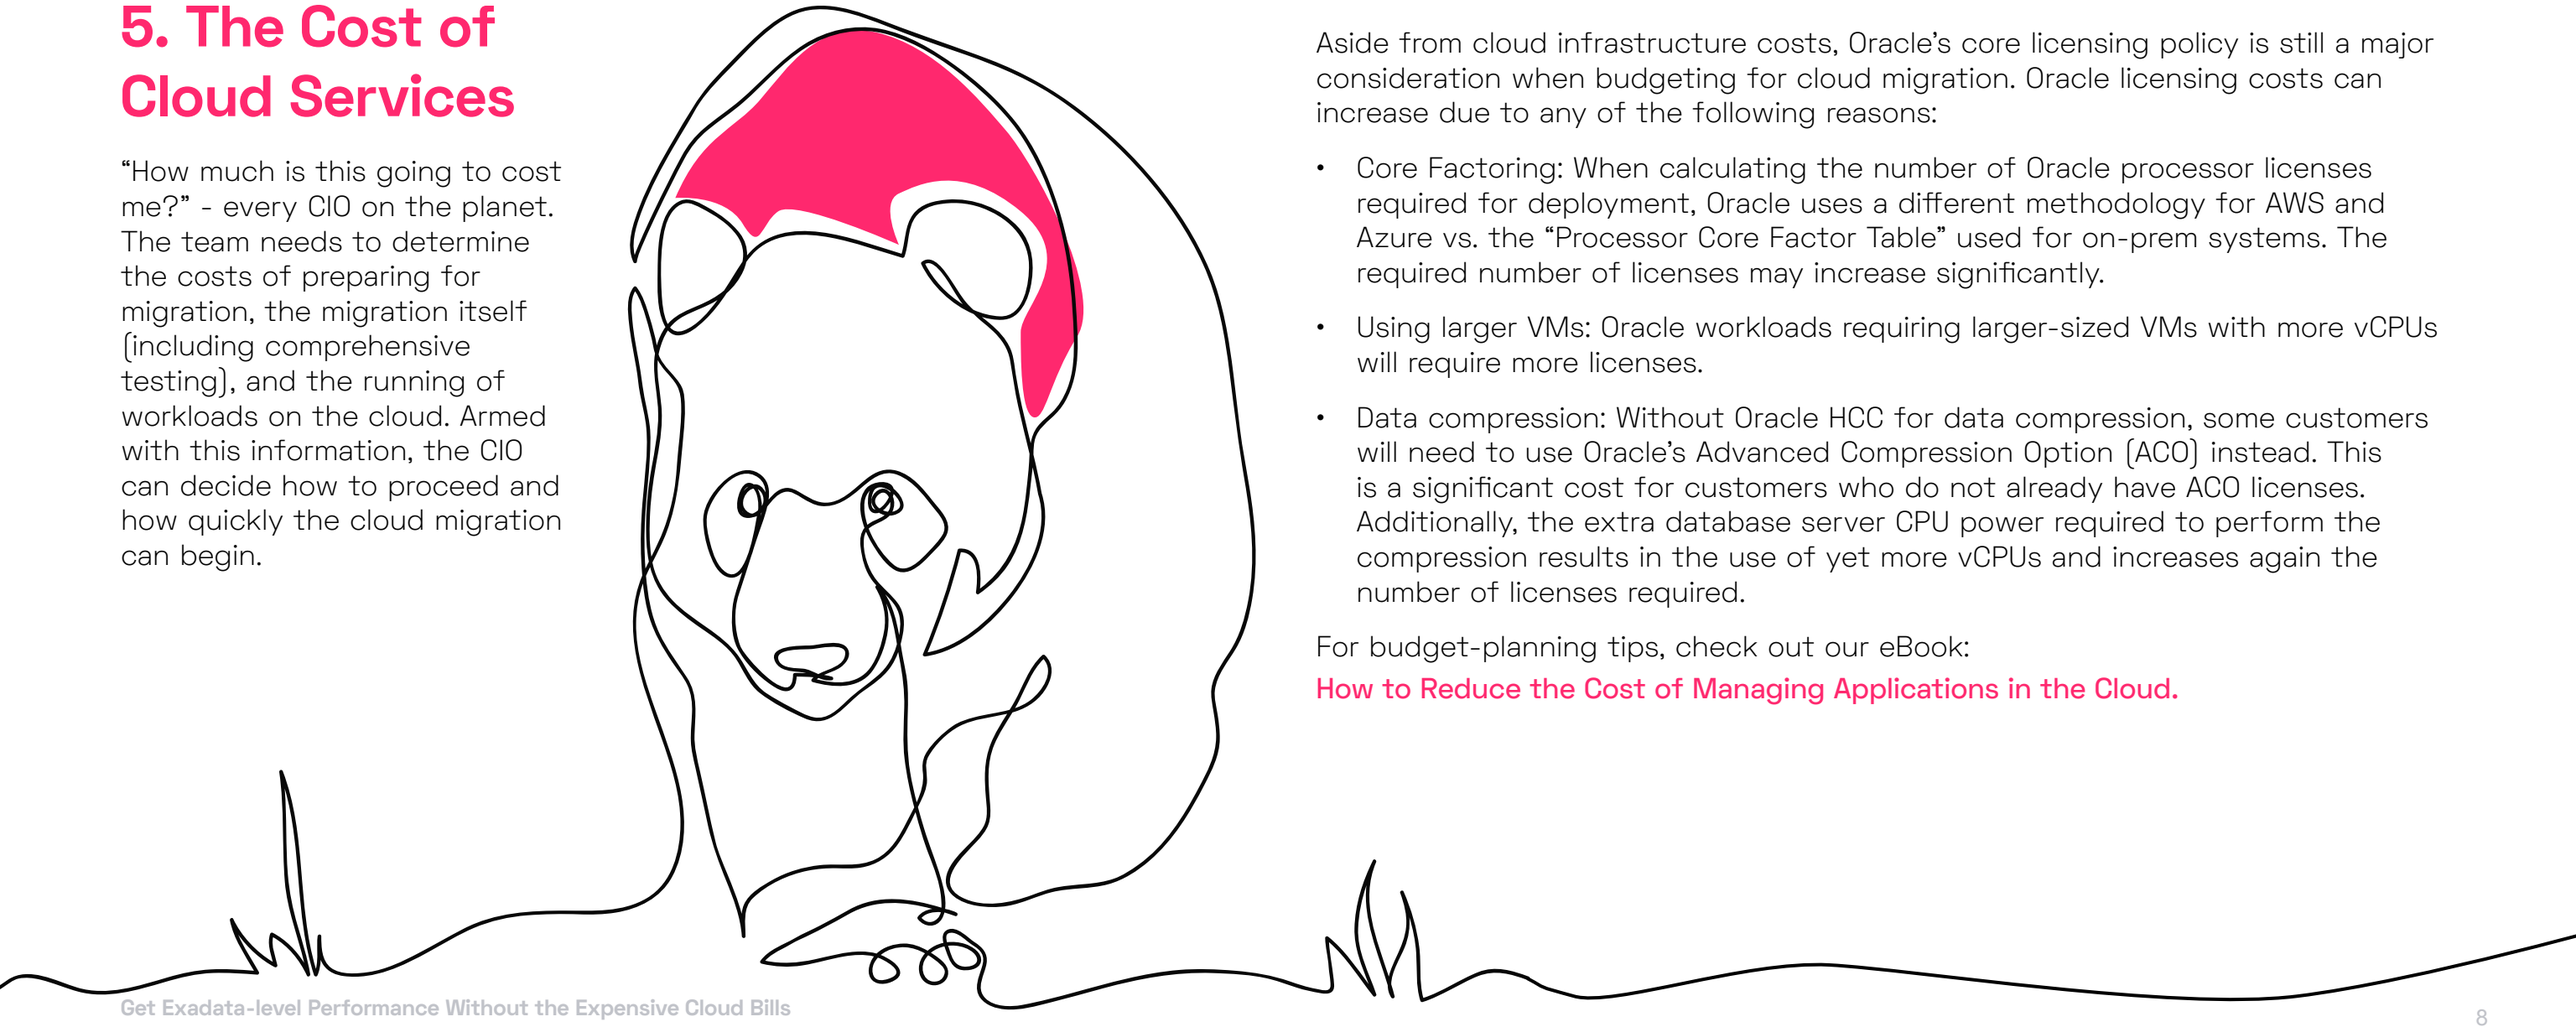

Aside from cloud infrastructure costs, Oracle's core licensing policy is still a major consideration when budgeting for cloud migration. Oracle licensing costs can increase due to any of the following reasons:

- Core Factoring: When calculating the number of Oracle processor licenses required for deployment, Oracle uses a different methodology for AWS and Azure vs. the "Processor Core Factor Table" used for on-prem systems. The required number of licenses may increase significantly.
- Using larger VMs: Oracle workloads requiring larger-sized VMs with more vCPUs will require more licenses.
- Data compression: Without Oracle HCC for data compression, some customers will need to use Oracle's Advanced Compression Option (ACO) instead. This is a significant cost for customers who do not already have ACO licenses. Additionally, the extra database server CPU power required to perform the compression results in the use of yet more vCPUs and increases again the number of licenses required.

For budget-planning tips, check out our eBook:
How to Reduce the Cost of Managing Applications in the Cloud.

Accelerate Cloud Adoption with Silk

Silk is the database supercharger, the smart data platform that sits invisibly between your Oracle databases and cloud infrastructure to deliver ultra-fast performance on the cloud. Silk's architecture makes cloud environments run up to 10x faster and the entire application stack more resilient to any infrastructure hiccups or malfunctions.

See How Silk Accelerates Cloud Adoption for Oracle Workloads

Full Data Mobility

Your data can move seamlessly across on-prem and cloud infrastructure to support hybrid and multi-cloud environments. Silk enables all data to flow freely as one.

No Refactor Required

Move workloads to the cloud without needing to change the application code or reengineer the dataset to work with a new database product. Silk enables easy lift-and-shift of Oracle databases to the cloud. You can still migrate data to PostgreSQL or to a managed PaaS offering down the road.

Faster Performance

Silk has proven performance of more than one million IOPS, 20 GB/s throughput, and consistent sub-millisecond latency, so it easily handles all OLTP, OLAP, and high-performance compute workloads on the cloud.

Scalability and Flexibility

Silk's unique scale up, in, and out architecture covers both dimensions of scale: capacity and performance. Silk linearly scales the performance by adding c.nodes (compute nodes) while independently scaling capacity by adding m.nodes (media nodes).

High Availability

Silk ensures data resiliency across zones, regions, and clouds. Silk data volumes are packaged up into data pods for easy migration to different infrastructures. If a failing resource is identified, a new one is preemptively provisioned from available cloud resources.

Instantaneous Zero-Footprint Clones and Data Replication

Quickly and easily create clones of Oracle database environments for unlimited development and testing. Zero-footprint clones have no impact on performance or cloud capacity.

Enterprise Data Services

Use thin provisioning, inline compression, and data deduplication to reduce your data footprint. These data services are often taken for granted on-prem but are not available natively on public cloud platforms.

Cost Efficiency

Silk reduces exposure to Oracle licensing costs and saves customers money on cloud bills. The architecture of Silk overcomes IaaS performance limitations, allowing the use of less VMs with less vCPUs that require Oracle licenses. Additionally, Silk offloads operations, such as data compression, to the data layer, further reducing the requirement for more vCPUs on the database server.

Database Consolidation

Silk's shared, multi-tenant capabilities mean that multiple, disparate databases can be consolidated onto one data platform, allowing you to achieve high performance for every workload while retaining cost-efficiency for all.

silk

Ready to see how Silk can help you achieve Exadata-level performance on the cloud?
Visit www.silk.us