

# VELOSTRATA

Overcoming the Challenges  
of Cloud Migration

WHITE PAPER

## Overcoming the Challenges of Cloud Migration

When considering public cloud, organizations need to start by evaluating which workloads will be a good fit. Prior to 2015, according to Gartner, enterprises primarily used the cloud for Mode 2 application workloads, which are characterized as dynamic in nature, developed in-house, follow an agile, continuous-integration model, and are able to accommodate downtime and errors. For these workloads, the main driver for migrating them to cloud has been agility. Since 2015, however, there has been a shift in the mindset of many CIOs. Enterprises are increasingly moving what Gartner would deem Mode 1 workloads to the cloud. Mode 1 workloads are typically static, IT-centric applications with substantial, persistent data that run major operations and need to adhere to high performance standards. Cost-efficiency is the main driver for this trend.

Although both types of workload present their own challenges, this white paper focuses on the considerations for migrating Mode 1 workloads to cloud.

There are several challenges that must be considered in order to mitigate migration risks.

### Addressing Application Complexity

In a recent survey conducted by Morgan Stanley, over 90% of CIOs reported that migration fell short of expectations, citing “complexity” as the main contributing factor. Indeed, migrating a complex, multi-tier Mode 1 application is a daunting task that requires careful consideration of multiple aspects, including:

- Complex system configurations with dependencies on external components/services, including external databases, file-shares, load balancers, subnet, and security.
- Bandwidth: Is there sufficient bandwidth on-premises allocated to the migrated application to support the application’s needs?
- Geography: Is the application chatty and latency sensitive? If so, will it work for users that are remote from the region in which the application service is deployed?
- Rightsizing: Since billing in cloud is per instance (vs. per-host billing on-premises), rightsizing can dramatically impact cloud costs.
- Image adaptation due to re-platforming, including driver injection, license configuration, and more.
- Cloud-based operations or post-migration configuration management, because on-premises tools are typically not suitable for cloud.

### Large and Live Data Migration

Applications that operate on large databases highlight the physical challenge of migrating large data-sets to cloud. For instance, moving 10TB over a 20Mbps link would take at least 50 days!

Furthermore, data migration is particularly problematic when dealing with transactional applications, i.e., applications that frequently change their data sets. For these applications, traditional replication-based approaches don’t usually work well. By the time replication of a data set has completed, the replica is already out of date and requires substantial synchronization. In addition, these applications might not be able to tolerate the downtime typically needed for cutover. Finally, these applications require backup, and sometimes high availability through synchronous replication, further exacerbating their migration to cloud.

### Vendor Lock-in

Enterprise customers are always concerned about being locked in to a particular vendor, because this implies they lose leverage in terms of cost, and they lose the choice to move to another (potentially more cost-effective) cloud, or shift workloads back on-premises. Most migration tools and cloud vendors lend themselves to a one-way migration process only.

### Security

Enterprise customers need to ensure that their migrated data is protected, both in transit to cloud as well as at rest. In addition, they have to ensure that their own on-premises data center is not compromised due to the migration. For instance, agent-based replication solutions install a replication agent on the workload machines on-premises, introducing a potential security threat.

## Cloud Migration Approaches

There are three main approaches to migrating workloads to cloud:

### Lift and Shift

Workloads are migrated to cloud IaaS “as-is” to the extent possible. In a pure Lift and Shift scenario, IT retains the same operations management tools and IT processes.

### Lift and Shift with Minor Config Changes

In a more advanced variant of Lift and Shift, IT managers move the workloads, but adopt some of the cloud-native operations and configuration management tools. They incrementally change IT processes and add cloud native tools such as auto-scaling and database as a service. Gartner refers to this advanced Lift and Shift as “Cloud-Enabled Virtual Automation.”

### Refactoring

In this mode, applications are re-coded, re-architected, and re-factored to use cloud-native services as much as possible. Typically, very few, if any, applications are migrated, and only the data needs to be migrated into cloud-native databases/repositories.

In this paper, we focus on the Lift and Shift approaches, which are the prevalent type of migration for Mode 1 workloads.

## Traditional Replication-Based Migration

The traditional migration process consists of the following stages:

### Discovery and Assessment

Discover application dependencies and assess which applications are a good fit for cloud and which are not. This step is carried out through a mix of manual and tool-based approaches, such as the AWS Discovery Service.

### Planning

Incorporate the output of a discovery tool into a “runbook,” consisting of application workloads, sequencing, and a scheduling plan.

### Migrate Workload

This is the phase where the migration is actually carried out, typically using a migration tool. Virtually all existing migration solutions are “replication-based.” That is, they essentially build a replica of the on-premises copy in the cloud and then switch to the cloud-based replica during the “cutover” period. In replication-based migration, this stage includes these six steps:

#### 1. Capture Application Environment

Ensure that the complete application is captured, including the application components, OS images, databases and configurations, as well as dependencies, runtime ordering, etc.

#### 2. Deploy Replica

In this phase, infrastructure is provisioned in the cloud, and the replica is actually deployed in the cloud. This involves the copying of all data and configuration. For large databases, this can take an extensive amount of time.

#### 3. Configure Replica

This step modifies the captured object to make it a replica that can co-exist without conflicting with the original app, including name-space changes (or isolated environments) and configurations, while ensuring the replica contains all of the necessary elements.

#### 4. Synchronize Replica

By the time the data has been moved to the cloud, data has changed on-premises, requiring additional synchronization. For transactional applications with heavy/frequent data changes, this is often another prolonged task.

#### 5. Test Workload in Cloud

Only now can IT managers get a sense of how the application will perform. If operations are not acceptable, the workload migration steps will need to be repeated.

#### 6. Cutover to Operation

This step involves downtime, reconfiguration, and restart.

## Velostrata: Streaming-Based Migration

Velostrata takes a unique approach to migration. Instead of replicating the application workloads, Velostrata moves, or “streams,” the actual application workloads to the cloud, thereby avoiding the sequence of steps involved in replication-based migration and significantly simplifying and speeding up the migration process. In addition, Velostrata provides a built-in recovery mechanism that allows consistent rollback of the application workload if the application misbehaves or the user wants to revert to the on-premises implementation for any reason. Finally, unlike replication-based approaches, Velostrata is agentless, thus mitigating security concerns and simplifying installation and configuration.

With Velostrata, once you have completed the Discovery and Assessment stages, the workload migration becomes a simple three-step process:

### 1. Clone-in-Cloud (test)

Using “snapshot-based” technology, an application runbook is cloned and then immediately run-in-cloud in an isolated environment (has no change in configuration). The workload is thin-provisioned on-premises, and similarly thin-provisioned in the cloud, i.e., the workload boots in the cloud accessing on-prem data through Velostrata’s caching technology. This allows detection of problems early in the process versus replication-based migrations whereby all data must be replicated before testing can begin.

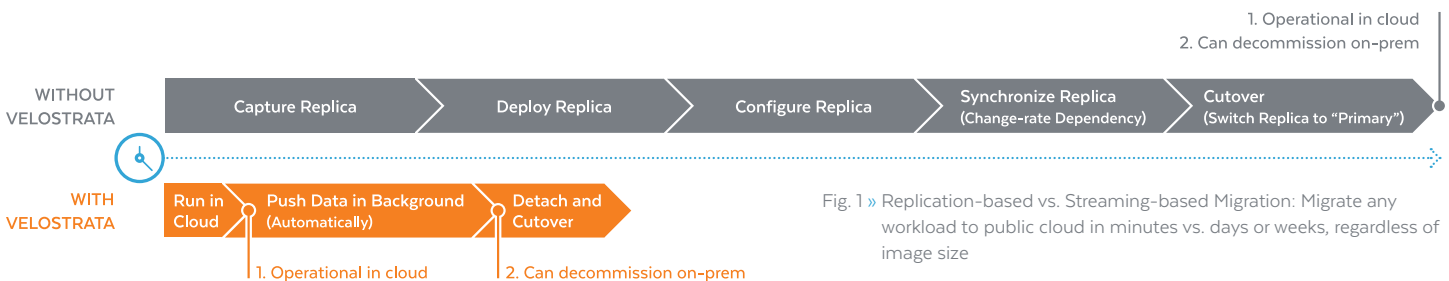
### 2. Run-in-Cloud (cutover)

Using the runbook for the application to be migrated, simply run the multi-tier apps in cloud. When invoked, each application shuts down on-premises and reboots in the public cloud in minutes, using Velostrata application streaming and boot over WAN technologies. At this point, the application is already operational in the cloud with good performance, thanks to the caching technologies. In other words, within a few minutes, cutover occurs. In the background, data is pushed to the cloud until fully cached. There is no need to synchronize data because there is no replica—all new data is written in the cloud. Nonetheless, Velostrata still maintains a fully synchronized copy on-premises for quick rollback purposes.

### 3. Detach

Data is moved from the caching system to native, block-based cloud storage and the system reboots as a native cloud application, detached from the rest of the system. At this point, on-premises resources can be de-commissioned.

The following figure illustrates the difference between the traditional replication-based flow and the Velostrata streaming-based flow:



## Velostrata: Core Technologies

To achieve this simple, cost-effective, and fast cloud migration, Velostrata leverages a number of innovative and patent-pending technologies.

A key architectural tenet, which enables workloads to be “cutover” and become operational in minutes within the cloud, is the ability to decouple compute from storage without sacrificing performance. Traditional data center design requires tight coupling of compute and storage. With Velostrata, compute instances can be deployed in the cloud in minutes while leaving the bulk data on-premises, accessing the remote data efficiently through a highly optimized, end-to-end channel that uses advanced WAN optimization technologies, such as multi-tier read/write caching, deduplication, compression, and pre-fetching capabilities.

A second key architectural innovation is the ability to boot-stream images over the WAN. Instead of replicating an image and converting it to cloud-native format (e.g., AMI), Velostrata boots the original on-premises image in the cloud over the network (network boot over the WAN), while adapting the image on-the-fly to the target cloud.

Since the instance is operational in the cloud right away, it conducts write operations that must be resiliently preserved while the instance is migrating to cloud. Furthermore, the changes must be absorbed in the cloud to avoid the latency penalty of writing the changes on-premises. To meet this challenge, Velostrata provides a highly resilient, multi-tier write cache that stores the updates in the cloud and is seamlessly merged with the original data-set as it is being migrated to cloud. Velostrata’s WAN optimizations include cross-VM deduplication, which allows migration of more VMs in parallel using the same available bandwidth when compared to replication-based solutions.

Data is written to two different availability zones and is further saved in an object store for high durability. In addition, Velostrata is crash-consistent, creating fresh, point-in-time snapshots on-premises to enable reverting from the cloud to on-premises operation if the application does not work as expected in cloud. Finally, Velostrata is agentless, which has several advantages. First, IT managers can migrate applications that are not currently running—saving time and resources. Also, there is no need to open network connections between the migrated applications and replication targets while they are running which can introduce security risk. Finally, Velostrata doesn’t require a lengthy downtime and cutover, and does not interfere with on-prem running apps.

The following figure summarizes the set of technologies developed by Velostrata:



Fig. 2 » Core technologies that differentiate Velostrata

## Summary: Velostrata Delivers Cloud Workload Mobility

Traditional replication-based migration is a complex and time-consuming process, especially for workloads involving transactional and frequently changing data, where consistent synchronization is time consuming and risky.

The streaming-based migration solution offered by Velostrata cuts by orders of magnitude the time by which applications become fully operationally in cloud—from days/weeks to minutes/hours. Furthermore, by moving the instance, not a replica, there is no need to conduct synchronization between replicas. Newly written content against the instance in cloud is saved in cloud, and unread content migrates in the background, while the application is already running in the cloud. Hence, cutover is guaranteed to be consistent, and fast. Finally, Velostrata also provides a built-in fast rollback to instantly revert applications back on-prem.

The ability to migrate production workloads in minutes while intelligently streaming data in the background enables true cloud workload mobility and supports key use cases such as cloud migration and dev/test. At the same time, you can reduce costs by migrating enterprise applications to the public cloud, thereby reducing operating expenses by simplifying and accelerating the migration process.

For more information about Velostrata architecture, visit [velostrata.com](http://velostrata.com) and download the “Understanding Velostrata Technology” white paper.