

Understanding the Velostrata technology for mass migrations into Google Cloud Platform



Table of Contents

Executive overview 3

Barriers to Cloud adoption 3

Overcoming adoption barriers 3

 Risk 3

 Complexity 4

 Cost 4

 Time 4

Accelerate mass migration 4

 Test before you migrate 4

 Compute in-cloud, storage on-prem 5

 Full migration to cloud 5

How Velostrata works 5

 Overview 5

 Deployment 6

 Streaming-based migration 7

 Resilient 9

 Simple and transparent management 9

 Safety and security 10

Summary and conclusion 10

Executive overview

This white paper will provide a deeper understanding of Velostrata’s unique technology and architecture, and will explore how these capabilities improve current mass migration practices. This paper is intended for a more technical audience whose responsibilities include data center infrastructure and/or public cloud.

More information about Velostrata is available on our website at velostrata.com.

Barriers to Cloud adoption

Public cloud infrastructure as a service (IaaS) promises agility, scalability, and a compelling “pay as you go” cost model. However, moving existing enterprise workloads to the cloud has always been challenging. Complexity abounds as companies struggle to adapt and migrate applications to run in a cloud environment. IT managers must understand application dependencies, change drivers and networking configurations, and learn new management interfaces.

Data gravity has further complicated cloud adoption and migration of production workloads—especially stateful applications with a lot of data. Traditional migration practices involve many labor hours, an assortment of tools, and substantial risk, not to mention the actual time required to move data from point A to point B. Moving lots of data to the cloud can also be expensive, as high performance block storage can be costly and you typically pay for it 24/7.

These barriers have prevented organizations from truly leveraging public cloud to their full capacity, either as an extension of their data center, where workloads can be moved dynamically to meet the needs of the business, or as a cloud-native data center. Velostrata helps organizations overcome these barriers—resulting in an agile IT organization that can be more responsive, effective, and cost-efficient.

Overcoming adoption barriers

Risk

It is not uncommon for organizations to end up with cloud migrations that have stalled or failed as a result of subpar pre-migration testing. That’s why Velostrata lets you test (and re-test) non-disruptively before you migrate, so you know exactly what to expect when your workloads arrive in the cloud for good.



Complexity

Adapting workloads to the cloud has required changing images, drivers, and storage. The process of replication, synchronization, and cutover has been complicated. Velostrata moves the actual application workloads (not a copy) and adapts them on the fly. This eliminates the need to copy and sync the data that changed after replication. Your applications run natively in the cloud in minutes and data is pushed to the cloud in the background.



Cost

Storage, especially high performance block storage, is usually the number one expense in the cloud. With Velostrata, you can reduce cloud storage costs by moving compute first and keeping storage on-prem until you're ready to completely cutover to the cloud. Plus, with Velostrata's fast time-to-cloud, IT can decommission double infrastructure on-prem faster, helping avoid costly hardware and software license refreshes sooner.



Time

Moving large amounts of storage takes a long time. For instance, moving 10 terabytes of data over a 20Mbps dedicated link takes 50 days to complete! To overcome this obstacle, Velostrata starts your applications in the cloud first, while storage migrates over time. As a result, your applications continue to run in the cloud with performance that feels local by leveraging Velostrata's advanced WAN optimization techniques. This means maintenance windows are short and predictable, and app performance is strong, both of which keeps IT and end users satisfied.

Accelerate mass migration

Velostrata delivers a fast, frictionless journey to the cloud for any organization no matter how big or small their IT team is. The Velostrata technology uniquely decouples compute from storage so you can run enterprise workloads in Google Cloud Platform (GCP) within minutes, while controlling and automating where storage resides. This unique architecture fundamentally changes how organizations execute their cloud strategies and migrations.

Test before you migrate

One of the many challenges of migration includes determining if an application will actually work in the cloud and how effectively. Velostrata makes it possible to move any workload in minutes and test application functionality and

efficiency. You can then compare cloud performance to on-prem and evaluate various cloud configurations (like instance type) to see what works best. Velostrata takes a snapshot of a live application with all its data and creates a test clone in the cloud. This lets IT test the application in the cloud non-disruptively, without any storage replication needed. Optionally, test data can be kept only in cloud, avoiding the need to set up extra storage on-prem or risk corrupting on-prem data.

Compute in-cloud, storage on-prem

Organizations may want to move compute to the cloud, but keep authoritative storage on-prem for security, operational, or cost reasons. Velostrata's unique technologies enable applications to run in the cloud without performance degradation, despite WAN latency, while data remains on-prem but synchronized with the cloud.

Full migration to Cloud

For organizations that want to migrate to GCP, Velostrata's streaming-based migration substantially reduces the time and complexity associated with traditional replication-based methods. Velostrata adapts the OS images on the fly, natively, for the target cloud and is up and running in minutes. Storage is efficiently moved while the applications are running, so there is no complicated re-sync and cutover, and no extended and unpredictable maintenance windows. Because a copy of storage remains on-prem until you disconnect it, you have a safety net to return a workload on-prem if needed.

How Velostrata works

Overview

Velostrata is a purpose-built, enterprise-grade cloud migration solution that lets you mass migrate enterprise workloads to Google Cloud Platform (GCP) from on-prem or other clouds quickly while controlling and automating where storage resides.

Our patented technology decouples compute (VMs) from storage (VMDKs) and provides intelligent streaming, optimization, multi-tier caching, and data pre-fetching capabilities to ensure optimal performance despite any WAN latency between on-prem storage and compute in the cloud. No manual changes to the applications, images, networks, or storage are required and IT organizations can leverage the same management tools and processes they use today.

With Velostrata, migrating production workloads to Google Cloud is as simple as a click of a button in our vCenter plug-in or a multi-VM sprint that's automated via our built-in runbooks. Velostrata also provides PowerShell and REST APIs for simple integration with third-party management solutions, as well as extensive monitoring capabilities.

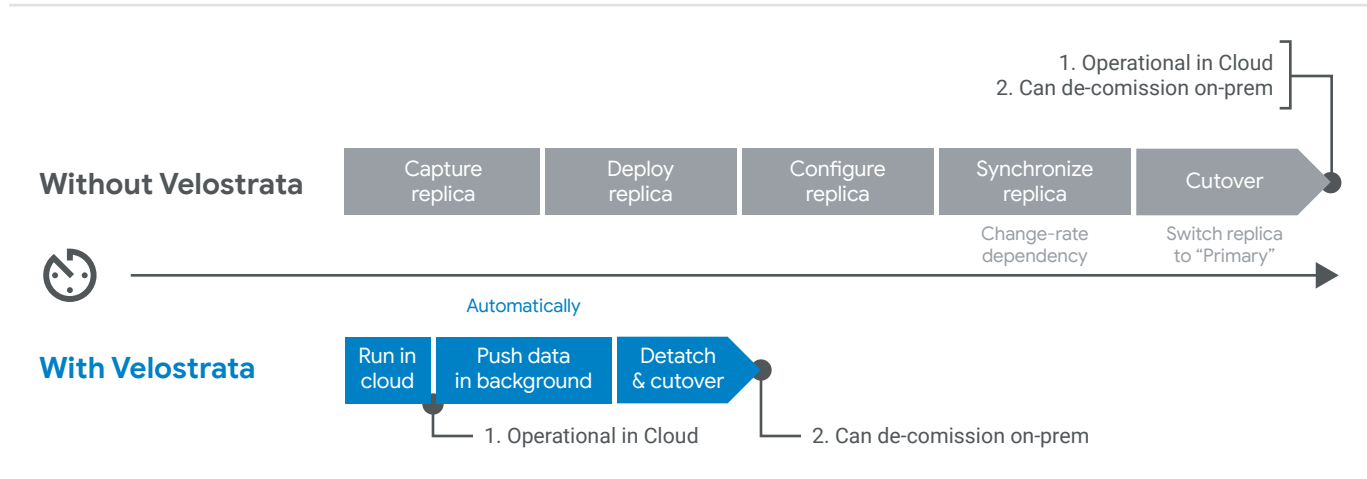
Deployment

Velostrata is a software-only solution, prepackaged and deployed in VMware as a virtual appliance — installation requires just a few easy steps. A Velostrata Management virtual appliance is deployed in your data center (or in GCP), and Velostrata Edge™ virtual appliances (called a Cloud Extension, or CE) are then deployed in a dual-node active/active configuration (across zones) in the target cloud(s) for scale and high availability.

Cloud Extensions are deployed within a customer- owned cloud account (VPC, for example) and connected to the on-prem virtualized infrastructure via a new or existing VPN connection. Velostrata encrypts data end- to-end, in flight and at rest. When using Velostrata in a hybrid model to run workloads in the cloud while keeping storage on-prem, each Edge virtual appliance supports up to 50 concurrent VMs.

More Cloud Extensions can be added as needed to scale out linearly. When using Velostrata for full migration, the breakdown is slightly different. A recommended configuration for a company that wants to migrate 1,000 virtual machines would be to deploy five Cloud Extensions to migrate 250 virtual machines in a batch, and the migration would be completed in four batches.

Figure 1. Accelerated migration with Velostrata



Velostrata can be managed via the Web Client or via the vCenter plug-in that integrates directly with vSphere.

Once Velostrata is installed and the cloud infrastructure has been set up, streaming workloads to and from the cloud is as simple as right-clicking VMs, selecting “run in cloud” and selecting a few additional deployment options. Alternatively, IT managers can use the PowerShell CLI and invoke a runbook script to automate migration for large numbers of multi-tier applications.

Velostrata’s unique, patented technology includes innovations in several key areas, including streaming- based migration, bi-directional optimization and multi- tier caching, resiliency, management, and security. These areas of innovation are covered in the sections below.

Streaming-based migration

Most solutions that migrate VMs to the cloud have many complex steps, including full replication of the images to the cloud, converting these images to the format used by Google Cloud, and then instantiating an image and booting it in the cloud. For transactional workloads that generate many data changes and can’t have significant downtime, tracking the data changes that occur during the migration and syncing them before the cutover to the cloud instance can be an additional challenge.

Velostrata differentiates from all other solutions by eliminating the need for replication. Instead, Velostrata performs a native boot of an on-prem operating system over the WAN in just a few minutes. While the OS boots, it is adapted on the fly by Velostrata to meet the requirements of the target cloud environment automatically and transparently, without any user intervention. Thus, unlike all other solutions, there is no need to replicate the boot image along with its entire virtual data disks as a prerequisite for booting the instance in the cloud.

Another major difference is that Velostrata is agentless. In contrast, most other migration solutions require customers to install an agent on each migrated server. Agent-based design leads to a number of undesirable outcomes.

First, the agent interferes with workload applications and competes with them on resources and opens up the door to a laborious cutover process that involves manually shutting down application services while the agent runs to complete a migration. Second, it opens up potential security holes by requiring an internet connection to the cloud target, possibly requiring A/V be disabled, and/or opening firewall ports to/from the application. Security risks aside, it adds tremendous manual labor overhead for any mass migration to the cloud.

Since with Velostrata the VM is not replicated, our technology streams only the data that is required for the workload to run and be available in the cloud. This approach is somewhat analogous to the way YouTube or other streaming video services work. Despite the length of the video, you can start watching instantly. If you fast forward, YouTube quickly downloads the new data without noticeable delay to the viewer.

In the case of VM disks, many parts may never be utilized at all. With Velostrata’s “run in cloud” operation, that data will remain on-prem unless queried. With Velostrata’s “full migration” operation, all data — even cold/unused data — will be migrated into the cloud.

Bi-directional optimization and caching

Velostrata includes a multi-tier, read-write cache in the cloud which combines in-memory (RAM) caching, SSD, and an object store.

For read purposes, the cache hosts the “working set” of data—this is the data that is frequently accessed by the application. Because the cloud VMs access this cache, they can achieve LAN-like performance. In addition, the cache includes sophisticated pre-fetching algorithms that predict data most likely to be needed next, further improving response times.

Velostrata also leverages its multi-tier caching for write operations, with latency that is lower than typical cloud-native block-level storage. In addition, Velostrata allows the user to choose whether changes should be written back to the on-prem storage, in which case the changes propagate asynchronously, or persist only in the cloud. This “cloud-persist” mode is useful for testing workloads in the cloud or for development environments.

Finally, Velostrata also provides strong bi-directional block level data de-duplication and compression algorithms to minimize the amount of data traversing the WAN and accelerate access to on-prem data. All written and migrated data is stored compressed, further improving storage cost savings.

The result is that for most enterprise applications, performance in cloud is on par with performance in the on-prem data center, despite the fact that storage and compute are now decoupled. In fact, in some cases performance actually improves with Velostrata’s approach because larger compute instances may be instantiated in the cloud and IOPs are now a shared resource that can scale out linearly, on-demand.

Resilient

All deployments of Velostrata include an Edge virtual appliance deployed in a dual-node, active/active configuration. Put simply, one Edge instance is deployed in one zone while a second Edge instance is deployed in a separate zone for redundancy and high availability. Velostrata acknowledges and ensures data resiliency by performing the write operations across these zones, so if one node fails, the system will continue to operate.

In addition, Velostrata stores the journal of write operations in Google's object store, Google Cloud Storage, to maintain a transient resilient data store while the data is written back to the virtual disks on-prem.

Velostrata keeps a maximum 30 seconds of write journal on the dual-zone nodes before that data is committed to the higher resiliency cloud object store. This 30 second RPO is relevant in the event that there are software issues (like a system crash) or an infrastructure issue (such as zone failure).

Velostrata's architecture ensures that there is never data loss related to a single Edge failure or data consistency issues. Further, Velostrata's architecture ensures a 30-second RPO for sync to cloud object store (resilient to dual-zone crash, which is very rare) and a 30-minute RPO for sync on-prem (resilient to whole cloud outage, which is extremely rare).

Simple and transparent management

With Velostrata, no changes to the applications, images, networks, storage, or drivers are required and there is no need to learn new tools/processes. Storage can be managed and backed up just as it is today, regardless of whether workloads have been moved to the cloud. In other words, your on-prem backup solution will back up all changes to VMs in the cloud, just as if they were on-prem. Velostrata also handles all image adaptation from the source (vSphere or other clouds) to GCP (and back) automatically, including injection of missing drivers, configuring the storage, licensing, and more.

Velostrata extends the actions of an existing VM object without replication or cloning, thus providing administrators with management context, continuity, and coherency. With Velostrata, there is no change to the managed object, no ambiguity, and no loss of administrative history or operational context. Management and reporting is integrated into the vCenter console through Velostrata's vCenter plug-in.

Velostrata is also designed with a PowerShell module and orchestrator for easy administration and automation of batch operations on bulk tasks, as well as a REST API for integration into third party management tools.

Migrating production workloads to the cloud with Velostrata is as fast and simple as right-clicking VMs and selecting “test in cloud”, “run in cloud”, “migrate to cloud”, or “run on-prem” (rollback).

Additional functionality includes selecting the instance type (larger instances or specific configurations may be selected to improve performance), setting storage policy (cloud persist or write-back), creating security groups, networking, and execution options, and more.

Safety and security

Security and data privacy is top of mind for most organizations today. Velostrata is configured to work within an existing or new VPN connection between the data center and the public cloud instance. Velostrata offers broad VPN support and will work with any VPN solution. Further, all data transferred by Velostrata is encrypted with AES-128 in-flight and AES-256 at-rest.

Summary and conclusion

Velostrata’s unique agentless streaming technology resolves all the previous barriers that enterprises were facing with cloud adoption and/or migration. With Velostrata, enterprises can mass migrate to the cloud with the speed, simplicity, and scale that they need to keep projects on track. And, since Velostrata eliminates agents and replication, it also vastly reduces risk and labor costs, while making planning easy to schedule and predict, thus eliminating a plethora of barriers that have hindered migration projects in the past. With Velostrata, the cloud journey that enterprises have been looking for has finally arrived.

For more information, please visit velostrata.com.

© 2018 Google LLC. All rights reserved. Google and the Google logo are trademarks of Google LLC. All other company and product names may be trademarks of the respective companies with which they are associated.

