

VELOSTRATA OVERVIEW

Table of Contents

Getting Started with Velostrata 3

 Introduction 4

 Terminology 6

 Technologies, Architecture, and Deployment 10

 Network Access Requirements..... 16

Getting Started with Velostrata

Introduction

Velostrata Summary

Velostrata gets enterprise applications running in the cloud within minutes while data migrates transparently in the background. Accelerate and simplify mass migration projects by supporting hundreds of migrations per day through Velostrata's patented streaming technology, agentless architecture, advanced WAN optimization, live in-cloud testing, analytics-based cloud instance rightsizing, and stateful rollback capabilities. With Velostrata, enterprises can validate, run, and migrate applications into Google Cloud Platform (GCP) without re-writing them, modifying the image, or changing management processes. Getting to the cloud is now as simple as a few clicks in vCenter. Velostrata does the rest.

Key Capabilities

- **Automatic and Seamless Adaptations** for cloud, including driver/agent installations, networking, licensing and more. No manual changes required to the applications, images, networks, storage, or drivers.
- **Provisioning and Rightsizing of Cloud Instances** using Velostrata, including support for customizations to networking, dedicated instances, encrypted disks, etc., as well as analytics-based rightsizing recommendations for both cost and performance.
- **Pre-migration Validation** via Velostrata's test-clone capability, which allows testing of production workloads and data directly in the cloud, within an isolated environment that has zero impact on production or live systems.
- **Migration Automation:** For pre-migration planning or migration automation, Velostrata offers self-documenting, auditable migration plans which IT can automatically generate based on vCenter Inventories. Automation runbooks provide a blueprint for each server, including their dependencies and configuration details, enabling scalable batch migration.
- **Stateful rollback** so no trip to the cloud is one-way if applications don't operate as planned.
- **Programmable Operation** with REST API or PowerShell enables automation of the migration flow using the orchestration/automation tools of your choice.
- **Customization Hooks** allow IT managers to perform specific customizations (e.g. enable or disable services) upon migration to cloud or rollback to on-prem.
- **Transparent Management:** No changes to the applications, images, networks, storage, or drivers are required and there is no need to learn new tools or processes as Velostrata integrates seamlessly into VMware vCenter. Velostrata is also designed with a PowerShell module, REST API for simple integration into 3rd party management tools.

These Guides

The upcoming sections will take you through the steps needed to deploy Velostrata and start migrating workloads into Google Cloud Platform (GCP). During this deployment guide, we'll start by defining the terminology and architecture that you'll see throughout the deployment. We will then outline the pre-requisites you'll need before you begin (along with instructions where appropriate to complete them). Finally, we'll proceed with the steps to properly install and configure your Velostrata deployment in the upcoming chapters.

If you need help along the way, you can always visit www.velostrata.com/support or email us at support@velostrata.com. We're here to help.

Terminology

If you're not familiar with Velostrata or GCP yet, it's important to understand the terminology that are going to be involved in this deployment. Let's define some of the key terminology you'll encounter throughout this guide:

Cloud Extension: An extension to customer virtual datacenter in the cloud. an environment in the cloud into which customer vSphere VMs can be migrated as part of the "run-in-cloud" operation. The Cloud Extension configuration includes:

- Selected cloud provider.
- Cloud region/location.
- Velostrata Role and User can be used for cloud API and resource access.
- Velostrata Cloud Edge Nodes configuration including subnets, security groups, availability zones/availability set.
- Configurations defaults to be applied on the workloads with the Run-in-Cloud wizard.

A Cloud Extension maps to a dual Cloud Edge Node setup.

Compute versus Storage: A Virtual Machine is defined by a set of associated resources of compute (for example, vCPU and RAM), storage as attached virtual disks and network interfaces. With Velostrata, a Virtual Machine can run in the cloud, where compute and networking resources are utilized from the cloud, while its virtual disks are kept in their original placement on-premises. Velostrata ensures storage consistency for the virtual disks on-premises as well as provides the expected storage performance required by the workload running in the cloud, even though the virtual disks are remote.

Velostrata Cloud Edge Node (A/B): The Velostrata Cloud Edge Virtual Machine is deployed in a dual-node configuration, where each node instance (A or B) is protecting the transient journal data (updates made to the data) of all VMs until it is persisted to the object store. The dual-node configuration exists in a highly available configuration (availability zones/set) which offers resiliency to cloud provider outages.

Velostrata Telemetry Service: Velostrata Cloud Edge and Velostrata Manager Virtual Machines report periodic performance and usage information to the Velostrata Telemetry Service, where this information is aggregated and processed at scale. The Velostrata Telemetry Service enables efficient monitoring of queries and activity graphing on-premises for Virtual Machines that are running in cloud.

Velostrata vSphere Web Client Plugin: A user-interface integration component, which registered with the vCenter Server to add Velostrata solution actions, operations summary and monitoring information to the vCenter user experience. Velostrata functionality is added to the vSphere Virtual Datacenter and Virtual Machine objects.

Velostrata Manager Virtual Appliance on vSphere: Includes the Velostrata Management Service and Datacenter Edge service. Deployed on vSphere where it has access to the vCenter Server, vSphere ESX hypervisors and the internet. The Velostrata Management Service is responsible for deployment, management and operations activity for Virtual Machines and Cloud Edge Virtual Machines.

Classless Inter-Domain Routing (CIDR): The set of IP addresses for the cloud.

Security Groups: Security groups act as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group. When deciding whether to allow traffic to reach an instance, all the rules are evaluated from all the security groups that are associated with the instance.

VPN: A Virtual Private Network (VPN) connection between your corporate datacenter and your cloud to create an extension of your corporate datacenter can be deployed on top of an internet connection or a DirectConnect/ExpressRoute connection.

GCP Specific Terminology

Regions and Zones: A region is a specific geographical location where you can run your resources. Each region has one or more zones. For example, the us-central1 region denotes a region in the Central United States that has zones us-central1-a, us-central1-b, us-central1-c, and us-central1-f.

VPC: Google Cloud Platform (GCP) Virtual Private Cloud (VPC) is a representation of your own network in the cloud. It provides networking functionality to Compute Engine virtual machine (VM) instances. It is a logical isolation of the Azure cloud dedicated to your subscription. You can fully control the IP address blocks, DNS settings, security policies, and route tables within this network. VPC provides global, scalable, flexible networking for your cloud-based services.

Cloud Deployment Manager: Deployment Manager is an infrastructure deployment service that automates the creation and management of Google Cloud Platform resources for you.

Firewall: GCP firewall rules protect your virtual machine (VM) instances from unapproved connections.

Network Tags: Network tags are used by networks to identify which instances are subject to certain firewall rules and network routes. For example, if you have several VM instances that are serving a large website, tag these instances with a shared word or term and then use that tag to apply a firewall rule that allows HTTP access to those instances.

Edge Network Tags: Velostrata Cloud Edge Components are deployed using a default dedicated edge network tag.

Workload Network Tags: Workload VMs are deployed using the default workload network tag.

Project: GCP projects form the basis for creating, enabling, and using all Cloud Platform services.

Service Account: A service account is the identity an instance or an application runs with. The Velostrata GCP deployment uses a management service account and a cloud edge service account.

For terminology related to other clouds which you may encounter, use the drop downs below.

AWS Specific Terminology

AWS Region: Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each region is a separate geographic area. Each region has multiple, isolated locations known as Availability Zones. Amazon EC2 provides the ability to place resources, such as instances and data in distinct locations. Resources aren't replicated across regions unless selected to do so specifically.

AWS Availability Zone: An Availability Zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Common points of failures like generators and cooling equipment are not shared across Availability Zones. Additionally, they are physically separate, such that even extremely uncommon disasters such as fires, tornados or flooding would only affect a single Availability Zone.

AWS VPC: Amazon Virtual Private Cloud (Amazon VPC) enables you to launch Amazon Web Services (AWS) resources into a virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of AWS. A Virtual Private Cloud (VPC) is a virtual network dedicated to your AWS account and is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC. You can configure your VPC; you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings.

AWS EC2: Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.

AWS EC2 Spot Instances: Amazon EC2 Spot instances allow you to bid on spare Amazon EC2 computing capacity. Since Spot instances are often available at a discount compared to On-Demand pricing, you can significantly reduce the cost of running your applications, grow your application's compute capacity and throughput for the same budget, and enable new types of cloud computing applications.

AWS IAM: AWS Identity and Access Management (IAM) is a Web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization).

Velostrata Security Group: The Velostrata Security Group in AWS/Azure includes the dual-node Velostrata Cloud Virtual Appliance. All Velostrata Cloud Edge Components are deployed into this dedicated security group that is separate from the security group for generic workloads.

Workload Security Group: The security group used for generic workloads.

Azure Specific Terminology

Azure Location: The Azure Platform is supported by a growing network of Microsoft-managed datacenters. The datacenters are grouped into what is referred to as locations. There are several different locations around the world.

Azure Resource Group: A resource group is a container that holds related resources for an application. The resource group could include all of the resources for an application, or only those resources that are logically grouped together. You can decide how you want to allocate resources to resource groups based on what makes the most sense for your organization.

Azure Availability Set: Microsoft periodically updates the underlying Azure fabric that is used to host VMs to patch security vulnerabilities and improve reliability and performance. These updates, which are referred to as planned maintenance events, are often performed without any impact on guest VMs. Sometimes, however, guest VMs must be rebooted to complete an update. To reduce the impact on guest VMs, the Azure fabric is divided into Fault Domains to ensure that not all guest VMs are rebooted at the same time. Each Application tier VM should be placed in an Availability Set. Velostrata uses an Availability Set to protect the availability of its core infrastructure VMs.

Azure VNet: An Azure virtual network (VNet) is a representation of your own network in the cloud. It is a logical isolation of the Azure cloud dedicated to your subscription. You can fully control the IP address blocks, DNS settings, security policies, and route tables within this network. You can also further segment your VNet into subnets and launch Azure IaaS virtual machines (VMs) and/or [Cloud services \(PaaS role instances\)](#). Additionally, you can connect the virtual network to your on-premises network using one of the [connectivity options](#) available in Azure. In essence, you can expand your network to Azure, with complete control of IP address blocks with the benefit of enterprise scale Azure provides.

Azure VM: Azure Virtual Machines lets you deploy a wide range of computing solutions in an agile way. Deploy a virtual machine nearly instantly, and pay by the minute. With support for Microsoft Windows, Linux, Microsoft SQL Server, Oracle, IBM, SAP, and Azure BizTalk Services, you can deploy any workload and any language on nearly any operating system. Provides persistent, durable storage volumes for use with virtual machines, and offers the option to select different underlying physical storage types and performance characteristics.

Azure AD/RBAC: Let's users securely control access to services and resources while offering data security and protection. Create and manage users and groups, and use permissions to allow and deny access to resources.

Velostrata Security Group: The Velostrata Security Group in AWS/Azure includes the dual-node Velostrata Cloud Virtual Appliance. All Velostrata Cloud Edge Components are deployed into this dedicated security group that is separate from the security group for generic workloads.

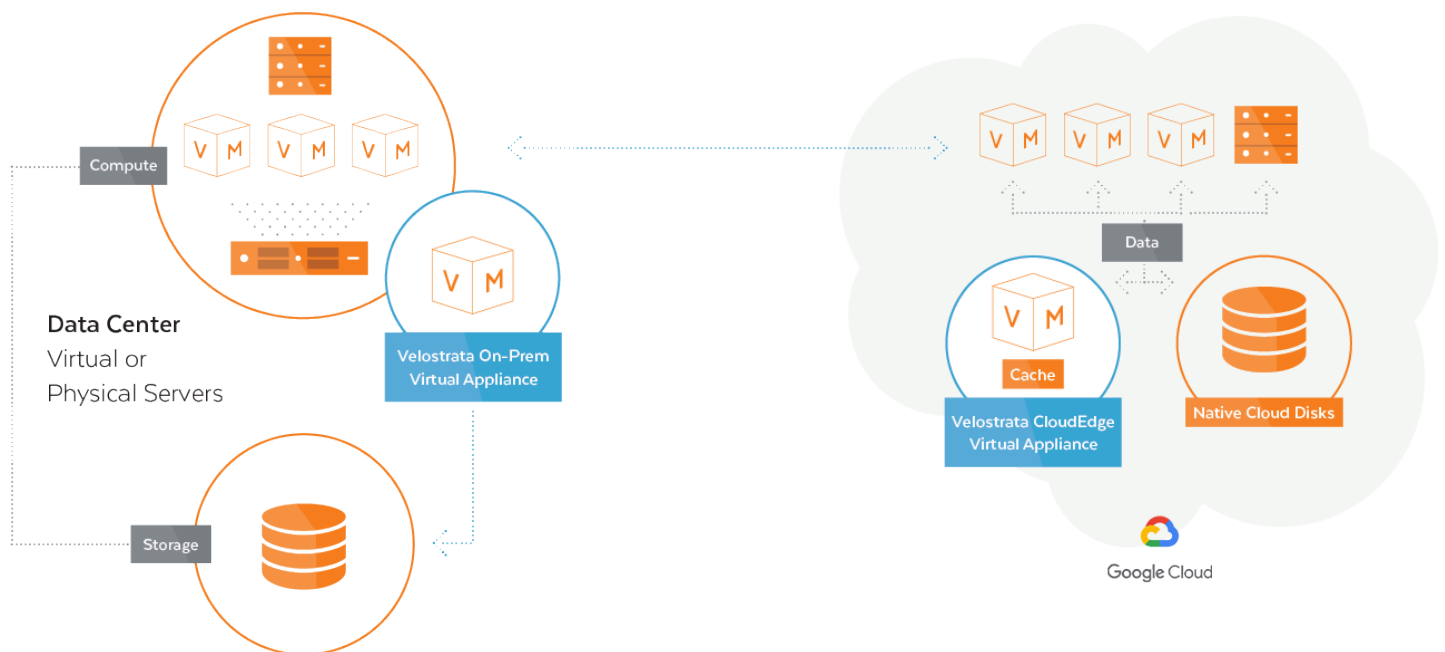
Workload Security Group: The security group used for generic workloads.

Technologies, Architecture, and Deployment

Architecture Overview

Velostrata's unique architecture decouples compute from storage in virtualized workloads, and introduces a number of patented capabilities and technologies that make cloud workload mobility possible. Our architecture was purpose-built to support enterprise-grade mass migration to the cloud:

- **Easy Deployment:** Quick, simple, and straight forward: Install the Velostrata virtual appliances in just a few steps and without installing agents on the servers you're migrating.
- **Simple Management in vCenter:** VMware vCenter plug-in reduces learning curve for VMware administrators. Migration operations are performed directly on your VM inventory (in vCenter) while tasks, events, and alarms provide full visibility and control over migration.
- **Secure By Design:** Data transfers between Velostrata virtual appliances use SSL and AES-128 encryption. Data at rest is de-duplicated, compressed, and encrypted with AES-256. Unlike SaaS-only migration tools, our appliance is fully controlled by the enterprise, which means cloud credentials are not shared with Velostrata.



Core Technologies

Velostrata's patented capabilities provide a frictionless path for enterprises to mass migrate their physical and virtual machines (VMs) from on-prem into the public cloud:

- **Boot Over WAN:** Velostrata performs a native boot from the on-premises operating system over the WAN, in just a few minutes— regardless of the image size. While the image boots, Velostrata adapts it on the fly to meet the target environment. No changes to the application, image, storage, drivers, or networking are required.

- **Intelligent Streaming:** Velostrata prioritizes necessary data for an application to run and moves that to the cloud first. Applications begin running in cloud within minutes, while less frequently accessed data can be streamed in the background while application runs in cloud.
- **Multi-Tier Caching and Optimization:** Velostrata includes a multi-tier, read-write cache in the cloud. This stores the working set of data needed by the application while it runs in the cloud. De-duplication, pre-fetching, asynchronous write-back and network optimizations further accelerate the migration and reducing migration bandwidth by up to 75% in production migrations.
- **Resiliency:** Each Velostrata Edge deployment incorporates active-active appliances deployed across two availability zones. Writes are acknowledged in both availability zones and then asynchronously transferred back to any on-prem storage to prevent data loss in the event of an outage. Optionally, new data writes can persist solely in the cloud, which is useful for dev/test scenarios. Velostrata's architecture ensures a 30-second RPO for sync in S3 and a 1-hour RPO for sync on-prem.
- **Supported Operating Systems:**
Windows Server: 2003, 2003 R2, 2008, 2008 R2, 2012, 2012 R2, 2016, 1709
Linux: RHEL 6.x, RHEL 7.x, CentOS 6.x, CentOS 7.x, SUSE Linux Enterprise Server 11 SP2 (or higher), SUSE Linux Enterprise Server 12, Ubuntu 12.x, Ubuntu 14.04.x, Ubuntu 16.x, Ubuntu 17.10.

Deployment Overview:

Velostrata software is deployed within virtual appliances. Once the networking and GCP prerequisites are met, installation takes place in a few simple steps:

1. Download and deploy the Velostrata Management Server (either on-prem, in GCP, or both - depending on your use case requirements).
2. Leverage Velostrata to create a Virtual Private Cloud, and a Virtual Private Network to securely link the Cloud provider's Virtual Network to your Velostrata Management Server(s). It's important to note that all traffic between these points is encrypted end-to-end, both in flight and at rest.
3. Deploy your Velostrata Cloud Extensions into the locations that you'll be migrating to and from (on-prem and GCP, for example). A Velostrata Cloud Extension supports the migration of 50 concurrent VMs, and you can add more Cloud Extensions for scale or to enjoy multi-cloud provider diversity.
4. During the migration, Velostrata provides comprehensive tracking and performance reporting, so you have a complete view of all your applications. In addition to pre-migration in-cloud testing and post-migration instant rollback capability options.

A typical Velostrata deployment consists of two parts:

1. Corporate Datacenter
2. Cloud(s) VPCs/VNETs

At the Corporate Datacenter:

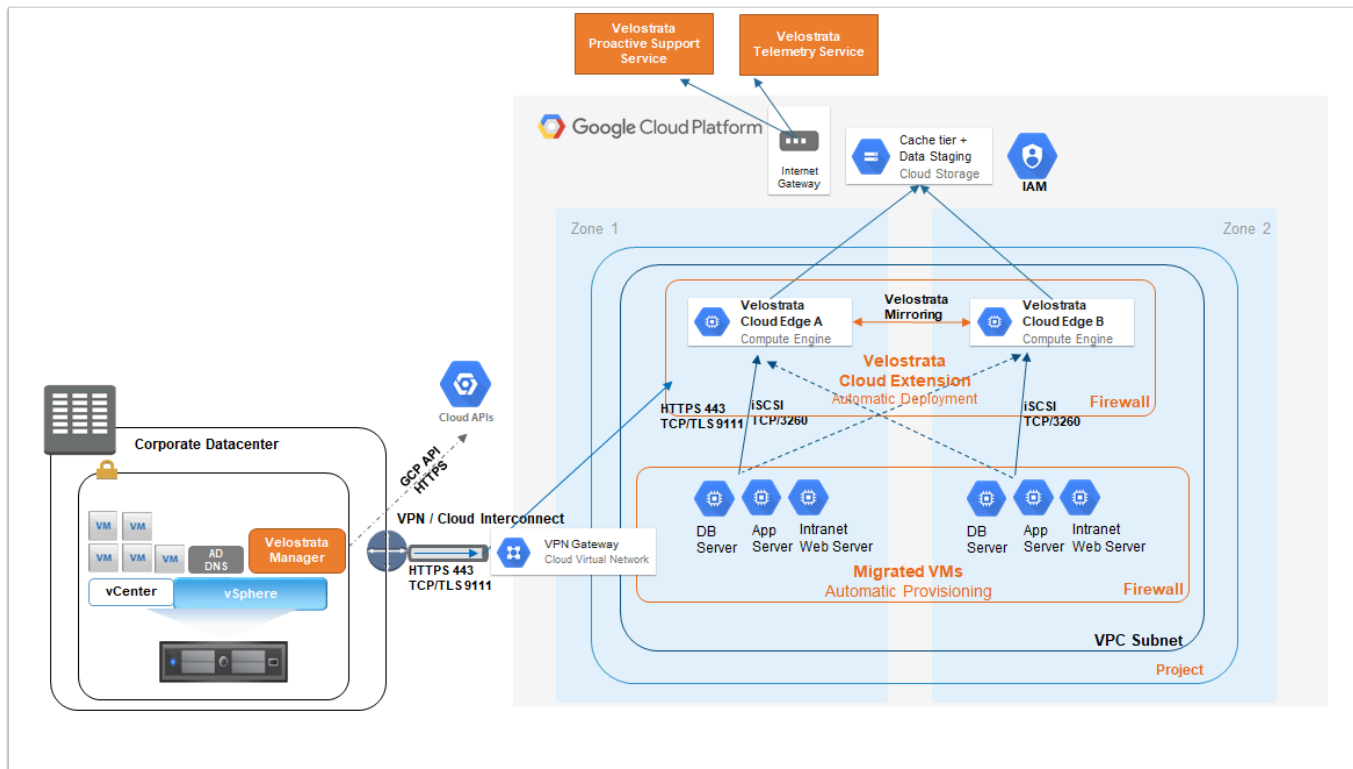
When performing on-prem to cloud migrations, the on-prem datacenter will be a crucial part of the Velostrata architecture.

- In the corporate datacenter, a VMware vSphere environment is deployed, running workload VMs. Infrastructure services such as DNS and Active Directory also exist in the datacenter. A Velostrata Manager Virtual Appliance is deployed in the vSphere setup.
- The Velostrata Manager Virtual Appliance will access the Cloud provider API endpoints on the public Internet as well as the Velostrata Telemetry Service, used for monitoring information and performance graphs when Virtual Machines run in cloud.
- The Velostrata Manager Virtual Appliance also hosts the Velostrata Datacenter Edge service, which is responsible for establishing the secure Datapath with the Velostrata Cloud Extension Edge Virtual Machine nodes in the cloud, as well as performing storage operations on Virtual Machine disks (VMDKs) attached to Virtual Machines running in cloud. The VMware Storage API is used for VMDK access.

In the cloud VPC/VNET:

- The Velostrata Cloud Edge nodes (A, B) are deployed automatically as part of the Velostrata Cloud Extension creation. The Velostrata Cloud Edge nodes are deployed into a separate security group to those used for workload Virtual Machines. This allows the creation of security policies that restrict access to and from the Velostrata components, the workload Virtual Machines and the datacenter. In general, only inbound access from the corporate datacenter is required for the Velostrata Datapath to operate, allowing for further tightening of security policies. Inbound access from workload Virtual Machines running in cloud is only required for storage connections to their virtual disks. These are inbound iSCSI connections from workload VMs into the Velostrata security group in cloud (connections to the Cloud Edge), not to on-premises.
- The Velostrata Cloud Edge nodes are in a highly available configuration, where two nodes exist in separate zones or regions. This ensures that during either a planned or unplanned maintenance event at least one Edge node will be available.
- The subnet(s) where the Cloud Edge nodes are deployed must allow access to the services, such as GCP Storage services (or Amazon S3 for cloud-to-cloud migrations). This allows access to both cloud services as well as outbound Internet access to the Velostrata Telemetry service. Inbound access from the public internet is blocked.
- Workload Virtual Machines (VMs) that run in cloud, whether using the Velostrata solution or whether launched directly using the cloud console, typically should run in private subnets that are not routed to the internet. These VMs should be routed to the corporate datacenter using the VPN solution, or access the internet using a NAT Gateway option.

GCP Deployment Overview



The diagram above depicts a typical Velostrata deployment. On the left hand side is the corporate datacenter (aka on-prem), and on the right hand side, the GCP VPC. The two are connected using a VPN or the GCP Cloud Interconnect option. It is also possible to have a deployment that adds an AWS cloud, in order to migrate VMs from AWS into GCP. It is also possible to have a cloud-to-cloud only architecture, in which there are no on-prem components, just AWS and GCP environments which communicate together (for migration out of AWS and into GCP).

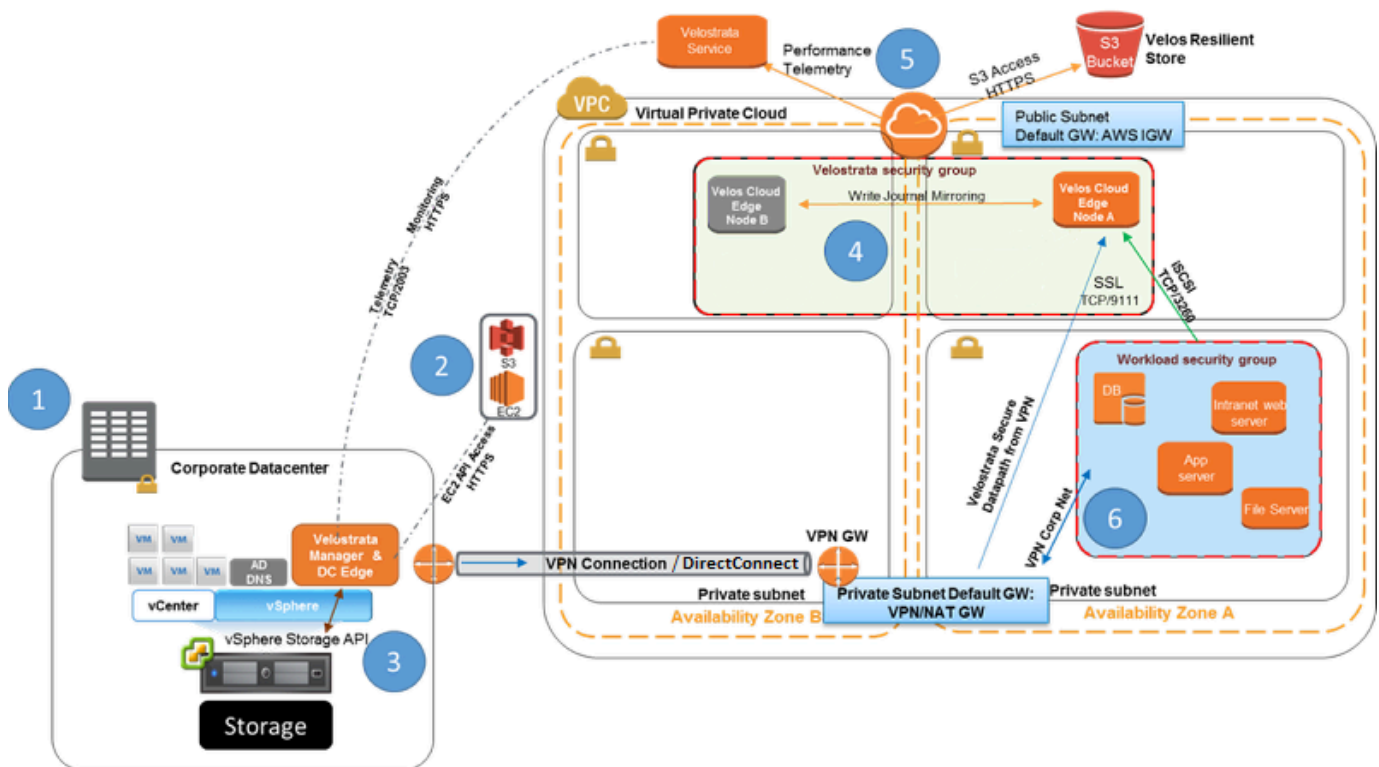
For more information on a recommended GCP VPC configuration, see [GCP Account and VPC Setup Requirements](#).

For more information on other clouds, please refer below:

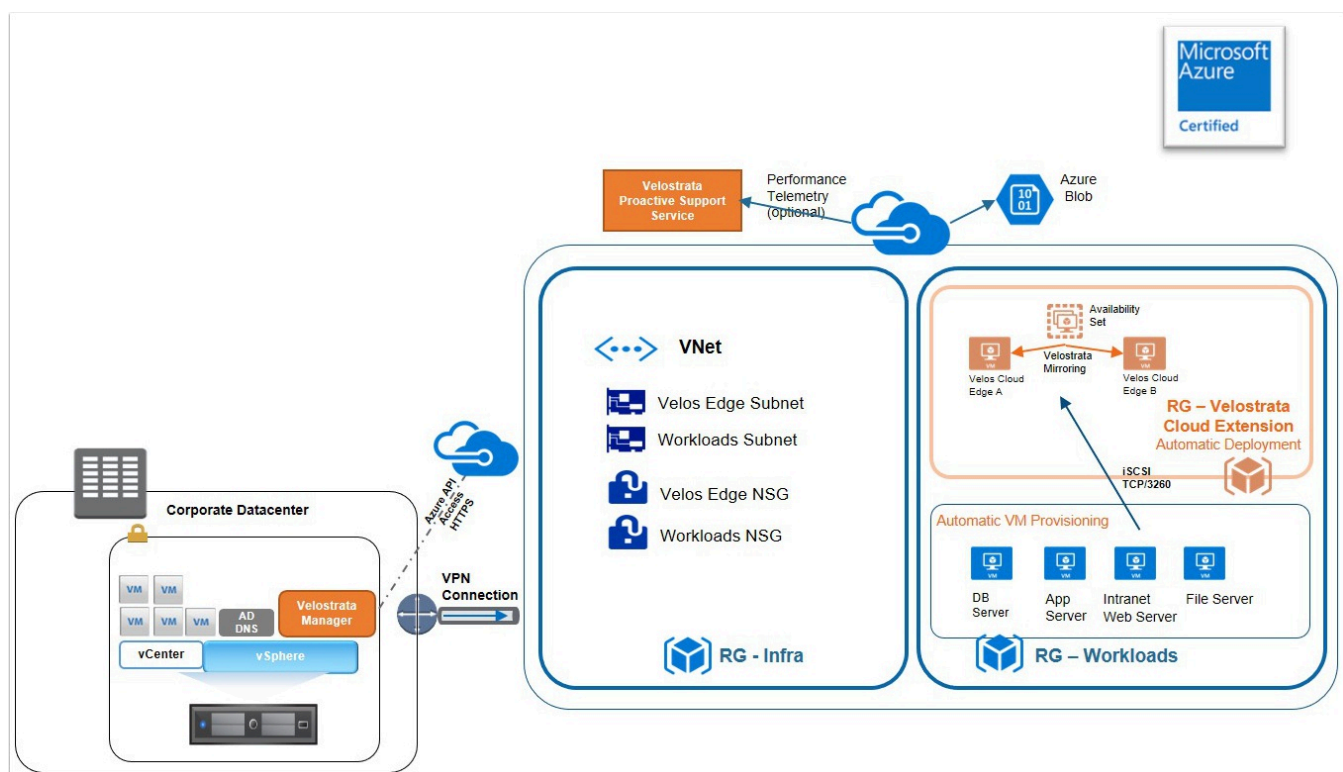
AWS Deployment Overview

The diagram above depicts a typical Velostrata deployment. On the left hand side is the corporate datacenter (aka on-prem), and on the right hand side, the Amazon VPC. The two are connected using a VPN or the AWS DirectConnect option.

For more information on a recommended Amazon VPC configuration, see [AWS Account and VPC Setup Requirements](#).



Azure Deployment Overview



The diagram above depicts a typical Velostrata deployment in Azure. On the left hand side is the corporate datacenter (aka on-premises), and on the right hand side, the Azure Resource Group, VNETs. The two are connected using a VPN or Azure ExpressRoute.

For more information on a recommended Azure VNet configuration, see [Azure Account and VNet Setup Requirements](#).

Network Access Requirements

In order for Velostrata to function properly, you will need to make sure your network, firewall, and VPN are all configured properly. Here is a table that explains the network access requirements that must be configured:

Source	Destination	Scope	Optional?	Protocol	Port
Velostrata Virtual Appliance on vSphere	vCenter Server	Corp LAN	-	HTTPS	TCP/443
	Velostrata Telemetry Service (optional)	Corp LAN	Yes	HTTPS	TCP/443
	vSphere ESXi	Corp LAN	-	VMW NBD	TCP/902
	Corp DNS Server	Corp LAN	-	DNS	TCP/UDP/53
	GCP/AWS/Azure API Endpoint	Corp Internet	-	HTTPS	TCP/443
	Velostrata Virtual Appliance in GCP/AWS/Azure	VPN	-	TLS/SSL	TCP/9111 TCP/443
	Workload Subnet in GCP/AWS/Azure – console probe	VPN	Yes	RDP	TCP/3389
		VPN	Yes	SSH	TCP/22
vCenter Server	Velostrata Virtual Appliance on vSphere	Corp LAN	-	HTTPS	TCP/443
Velostrata Edge Network Tags (GCP) [Security Group (Azure/AWS)]	GCP Storage [AWS S3 Endpoint/Azure Storage]	AWS/Azure/ GCP Internet	-	HTTPS	TCP/443
	Velostrata Telemetry Service (optional)	AWS/Azure/ GCP	Yes	HTTP or HTTPS	TCP/443

Source	Destination	Scope	Optional?	Protocol	Port
		Internet			
Workload Network Tags (GCP) [Workload Security Group (Azure/AWS)]	Edge Network Tags in GCP [Velostrata Security Group in AWS/Azure]	AWS VPC/AZURE VNET/GCP VPC	-	iSCSI	TCP/3260
			Yes	SYSLOG (for boot phase)	UDP/ 514 (optional)
	Workload dependent	Corporate LAN	Workload dependent	Workload dependent	Workload dependent

Continuing Deployment...

All done? Continue to the next guide, [Velostrata Deployment](#).