tyco | American Dynamics

victor Video Management System

Cloud Deployment Support Guide

A resource for IT practitioners planning to deploy and support the victor Video Management System (VMS) on Cloud Infrastructure as a Service (IaaS)

Published November 2019
Document Version 1.0

Johnson Controls

1 Table of Contents

2 About this guide

This document is provided “as-is” for informational purposes only. The topics and references in this document are subject to change without notice.

This document is intended for IT practitioners familiar with victor VMS architecture, Infrastructure as a Service (IaaS) and the Shared Responsibility model for Cloud security.

All guidance and references in this document are limited to Cloud Service Provider (CSP) offerings from Microsoft Azure, Amazon Web Services, and Google Cloud Platform.

For purposes of simplicity and the underlying technology of the victor VMS application, many diagrams in this document will use Microsoft reference illustrations to support general cloud concepts that may equally apply or translate across CSP architectures.

Johnson Controls makes no guarantees or warranties, expressed or implied, as to the supportability of its application based on information in this document.

Document information:

Document title: victor Video Management System Cloud Deployment Support Guide

Document version date: 18 November 2019

3 Introduction

As private and public cloud adoption continues to grow, on-premise lift and shift strategies are accelerating and have become a top focus for many organizations. Corporate IT teams are taking a stronger role in cloud adoption with focus on infrastructure cost management, security governance, and advising on critical workloads that can and should move to cloud.

This guide will provide general recommendations to assist customers with successful planning and deployment support for hosting their victor VMS on Infrastructure as a Service (IaaS).

All guidance provided in this document is based on victor application server v5.4.1 release or later running natively on IaaS provisioned virtual machines for Microsoft Windows Server 2016 and Microsoft SQL Server 2016 database. ***Currently, victor does not support SQL Relational Database as a Service (DBaaS) options from any cloud service provider and should not be considered in system design for implementation.***

The victor VMS is architected as a 3-tier client-server application, supporting Standalone single server, victor Professional VMS, and Enterprise multi-server system deployment models, victor Enterprise VMS, for different operational needs. Regardless of the system deployment model, a fundamental principal that should be understood and considered in system design is that any victor application server can be configured as a single or split application and database server:

1. Single All-in-One Server (victor application and SQL DB on same server).

2. Split Application & DB Server (victor application server pointing to a remote SQL DB instance).

The general approach for the actual installation and support of victor on cloud IaaS is not significantly different to on-premise virtualized deployments using Hyper-V or VMWare solutions. Although in contrast, one of the main benefits of using cloud IaaS is the ability to more accurately monitor and provision specific resources as you grow, without needing to over provision for maximum capacity upfront when it is not required.

Cloud IaaS provides tremendous flexibility and configuration options requiring careful consideration for a successful victor deployment that is optimized to meet the different requirements and constraints of each customer environment and varying workload demands.

4 Shared Responsibility Model for Cloud Security

The roles and responsibilities for complete security in the cloud are not significantly different than on-premise strategies and practices, most of the security responsibility largely remains with the customer. The shared responsibility model commonly applies across all CSPs relative to the infrastructure service offering and basically separates roles and responsibilities into two categories:

1. **CSP owns "Security OF the Cloud"** typically limited to the physical security of data centers and the infrastructure inside such as Compute, Storage, and Network.

2. **Customer owns "Security IN the Cloud"** which is all information security related to access control, data, and protection of infrastructure consumed as a service. For IaaS, customers have full ownership and control of all their data, applications, operating system and network configuration, so they are responsible for securing it.

Figure below extracted from Shared Responsibilities for Cloud Computing; Microsoft published April 2017.

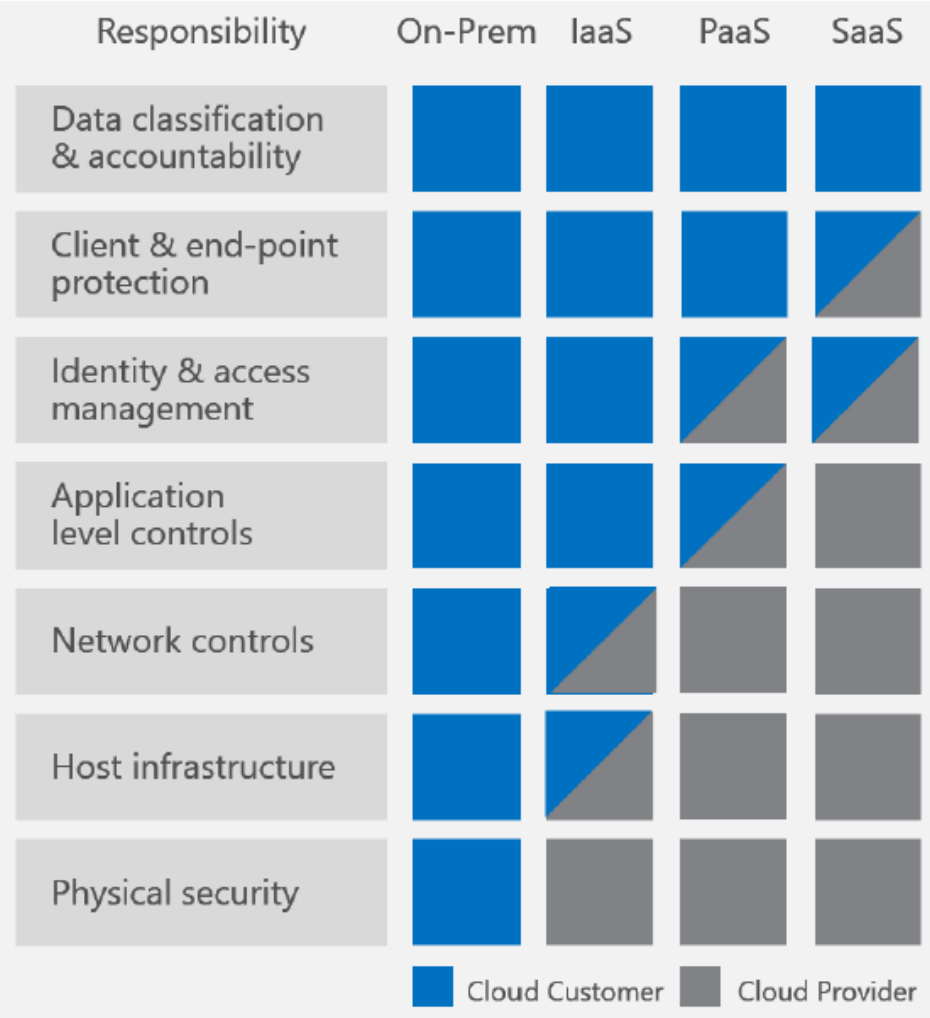

For information resources specific to CSP cloud security and the shared responsibility model, you can reference the site links for **Microsoft Azure**, **Amazon AWS**, and **Google GCP**.

5 Premise to Cloud Network Connectivity

Customer specific IT infrastructure and security strategies will drive the fundamental “design for failure” requirements to securely extend on-premise corporate networks and remote sites to virtual private cloud environments based on CSP specific architecture and service offerings.

A victor application and database server can be securely hosted on VMs in cloud IaaS, providing it can adequately service client connectivity and workload bandwidth demands that predominantly originate on-premise and in many cases are highly distributed across multiple sites.

victor Client connections include:

1. victor Thick client applications
2. victor Web client applications
3. victor Mobile client applications
4. victor integrations for on-premise applications and devices
5. VideoEdge NVRs for physical access, video and I/O management

When considering the deployment and operation of victor over a hybrid premise to cloud network, VPN and Dedicated Network service options are the most critical to consider in “design for failure” strategies.

Figure below represents typical customer premise connectivity options to an isolated and private network environment in cloud IaaS (Azure architecture reference example).

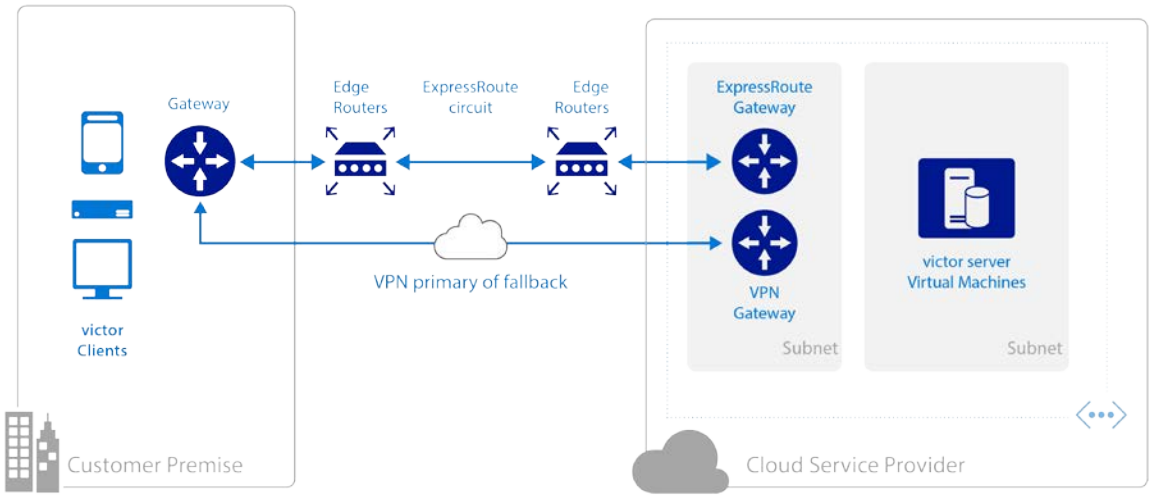

VPN is simply a secure way to extend and connect on-premise networks to a CSP network through IPsec VPN tunnels, traffic is encrypted and traverses the public Internet between the two “networks. In general VPN is a viable option as primary connectivity for low volume traffic or as a fall-back redundant path to a higher speed link. The throughput of the on-premise gateway has impact if not properly configured or able to support the equivalent or higher VPN service throughput capacity.

When VPN may not provide the required performance, security, or connectivity typology needs of large complex networks, Dedicated Network connectivity via fiber transport is typical with service models for (1) Customer Managed Interconnect through a Colocation/Exchange Provider and (2) Carrier Managed Interconnect through an Intermediate/Transit Carrier.

The following table serves as a comparison reference for all Network Service offerings by CSP to understand and consider in design.

Network Service Type	Azure	AWS	GCP
Cloud Virtual Network - Isolated and private customer environment in the cloud.	Virtual Network	Virtual Private Cloud	Virtual Private Cloud
Dedicated Network - dedicated and private network connection from a customer location to cloud provider.	ExpressRoute	Direct Connect	Cloud Interconnect
Cross Premises Connectivity - Connects virtual networks to other virtual networks or customer on-premise networks.	VPN Gateway	API Gateway	Cloud VPN
DNS - Domain name system management and user to application availability routing.	Azure DNS Traffic Manager	Route 53	Google Cloud DNS
Load balancing - Automatically distribute incoming application traffic to add scale, handle failover, and route to resources.	Load Balancer Application Gateway	Elastic Load Balancing	Cloud Load Balancing

6 Business Continuity & Disaster Recovery

Data protection and system availability objectives are typically driven by compliance requirements and cost to manage acceptable level of risk… in other words, what can you afford to lose when something bad happens? Several key terms, including Recovery Time Objective (RTO) and Recovery Point Objective (RPO), have emerged to help define business requirements and measure how well solutions can satisfy these requirements.

RPO refers to the amount of data measured in time, which is tolerable to lose when a disruptive event causes system data loss or corruption. RPO largely determines the frequency of data backups or snapshot replication required.

RTO refers to how much time is tolerable to lose after a disruptive event causes a system outage, and it returns into service. Generally, this relates to the entire time it takes to operationalize. RTO largely determines the type of system redundancy, failover, and recovery orchestration required to restore system availability.

There are multiple components to consider in total system deployment design for comprehensive business continuity to address planned and unplanned service-interrupting events that impact application availability. The two most fundamental and critical components to consider are Storage offerings for data protection (data replication, backup, restore) and VM instance application protection (machine replication, standby, failover).

Figure below is a high level view of the Business Continuity & Disaster Recovery continuum in Cloud infrastructure that can be achieved, extracted from [Azure Resiliency infographic](), Microsoft published March 2018.

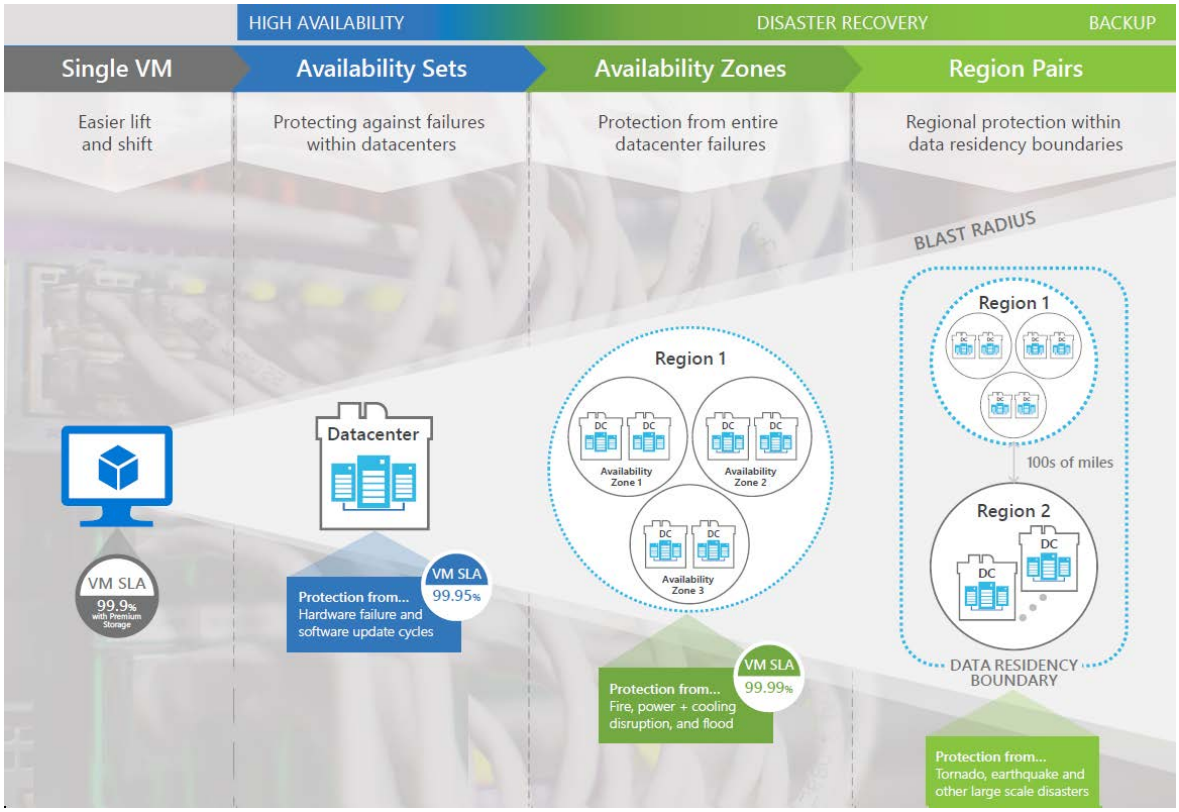

Although each CSP approaches business continuity and disaster recovery options differently. Each with Pros/Cons in degrees of setup complexity, configuration flexibility, costs, as well as 3rd party packaged solutions support. In general, Microsoft Azure provides the most native out-of-the box business continuity and disaster recovery options for ease-of-use with Microsoft application infrastructure advantages, while Amazon and Google take more of a design and build your own approach based on individual service offerings.

The following section provides a brief summary with reference links to better understand CSP specific approaches and options to consider for cloud architecture choices and system deployment design.

6.1 Microsoft Azure

Azure Backup Service – backs up VMs to a Recovery Service Vault. Daily/Weekly/Monthly scheduling with retention policies for snapshots. Individually recover data files and folders or restore a VM from saved recovery points.

Azure Site Recovery Service (ASR) – replicates VMs from primary region to secondary region for manual DR fail-over and failback during planned and unplanned outages.

- Crash-consistent recovery points every 5 minutes
- App-consistent snapshot frequency configurable from 1 to 12 hours
- Recovery point retention configurable from 1 to 72 hours

Azure VM Storage Spaces Direct (S2D) – is software defined storage which provides a way to create guest clusters on Azure. A guest cluster in Microsoft Azure is a Failover Cluster comprised of IaaS VMs. It allows hosted VM workloads to fail over across the guest clusters achieving higher availability SLA for applications than a single Azure VM can provide.

Azure VM support options for SQL Server Always On -– Availability Groups and Failover Cluster Instances.

6.2 Amazon AWS

Amazon provides **white paper** guidance for backup and recovery approaches to building business continuity and disaster recovery solutions based on AWS Backup & Restore Services.

AWS Backup & Restore services – provides the list of individual storage services, data-transfer methods, and networking options for data protection and high availability.

AWS CloudEndure Disaster Recovery – an AWS service to shift DR strategies to the AWS cloud from existing physical or virtual data centers, private clouds, or other public clouds. For customers having already migrated to AWS, they can further protect mission-critical workloads with cross-region disaster recovery.

EC2 VM support options for SQL Server Always On – Availability Groups and Failover Cluster Instances.

6.3 Google GCP

Google provides **white paper** guidance for using GCP products and services as building blocks to design application architectures for business continuity and disaster recovery needs. GCP architectures to consider:

1. **Cold: recoverable application server**
2. **Warm: Static Site Failover**
3. **Hot: HA web application**

Compute Engine support for SQL Server Always On – Availability Groups and Failover Cluster Instances.

6.4 victor System Considerations

In summary, although CSP specific products and services will vary and influence architecture choices, the goal is to design and deploy the underlying Windows Server and SQL DB Server infrastructure with desired business continuity and disaster recovery objectives that the victor application will depend on to run.

Once in cloud, each CSP provides plenty of tools for monitoring infrastructure and virtual machine health, but most critical is victor application state monitoring to facilitate manual or automated application recovery procedures to meet RTO/RPO targets. Two victor application instances can be individually licensed, deployed, and managed for business continuity needs providing only one instance is active at a time, while the other remains in passive standby. A Hardware Flexibility License (HWFLX) is required per victor license unless the second license is used solely for redundancy to mirror the databases only.

7 Sizing Compute, Memory, Storage

A victor application server with a MS SQL Express database server can be installed and will run on a single Windows Server 2016 OS virtual machine configured as small as (2 CPU, 4 GB memory, 128 GB Standard SSD single disk), and may be suitable for non-critical light use applications including Dev/Test scenarios.

When sizing a VM for production critical workloads, there are a multitude of design considerations to optimize cost and performance for customer specific victor application demands. Unlike traditional on-premise infrastructure, IaaS provides significant benefits in the ability to more easily monitor VM resource consumption relative to sustained performance impacts (CPU, Memory, Storage IOPS) from application demands so one can reconfigure or change a VM size when required. A rule of thumb is not to over provision VM resources upfront since you can more easily provision up when needed, since having to de-provision compute resources typically requires creating and deploying a new VM.

When considering a victor server deployment, VM size configurations will vary by CSP but the following provides minimum recommendations:

- General Purpose, Compute Optimized, or Memory Optimized VMs with Persistent Disk Storage support are all appropriate starting points for victor workloads.
- Min 4 CPU, ~16 GB RAM, 128 GiB SSD OS boot disk with Windows Server 2016
- Secondary 128 GiB SSD data disk for victor and/or MS SQL Server 2016
- For a typical single VM deployment (All-in-One victor application & DB Server), MS SQL Server 2016 Express or Standard can be considered and resourced accordingly based on transaction processing demand.
- For a multi VM deployment (Split victor Application & MS SQL DB Server), MS SQL Server AlwaysOn Availability Groups and Failover Cluster Instances should be considered to address all aspects of database HA & DR.

7.1 Microsoft Azure VM minimum recommendations

Azure D-Series General Purpose D4s v3 with Premium SSD Managed Disks.

- 4 CPU, 16 GB RAM, 128 GiB SSD OS boot disk with Windows Server 2016
- Secondary 128 GiB SSD data disk for victor and/or MS SQL Server 2016
- SSD disks can be provisioned up to accommodate specific storage capacity and/or increased IOPS throughput performance needs.

Microsoft Azure site links for **Windows Virtual Machines** and **Managed Disk** options.

7.2 Amazon AWS VM minimum recommendations

Amazon EC2 General Purpose m5.xlarge with EBS SSD Disks.

- 4 CPU, 16 GB RAM, 128 GiB SSD OS boot disk with Windows Server 2016
- Secondary 128 GiB SSD data disk for victor and/or MS SQL Server 2016
- SSD disks can be provisioned up to accommodate specific storage capacity and/or increased IOPS throughput performance needs.

Amazon AWS site links for **EC2 Instance Types** and **EBS Storage** options.

7.3 Google GCP VM minimum recommendations

Google Standard Machine Type n1-standard-4 with SSD Persistent Disk Storage.

- 4 CPU, 15 GB RAM, 128 GiB SSD OS boot disk with Windows Server 2016
- Secondary 128 GiB SSD data disk for victor and/or MS SQL Server 2016
- SSD disks can be sized up to accommodate specific storage capacity and/or increased IOPS throughput performance needs.

Google GCP site links for **Compute Engine Machine Types** and **Storage** options.

Figure below represents a typical single VM minimum configuration (Azure architecture reference example).

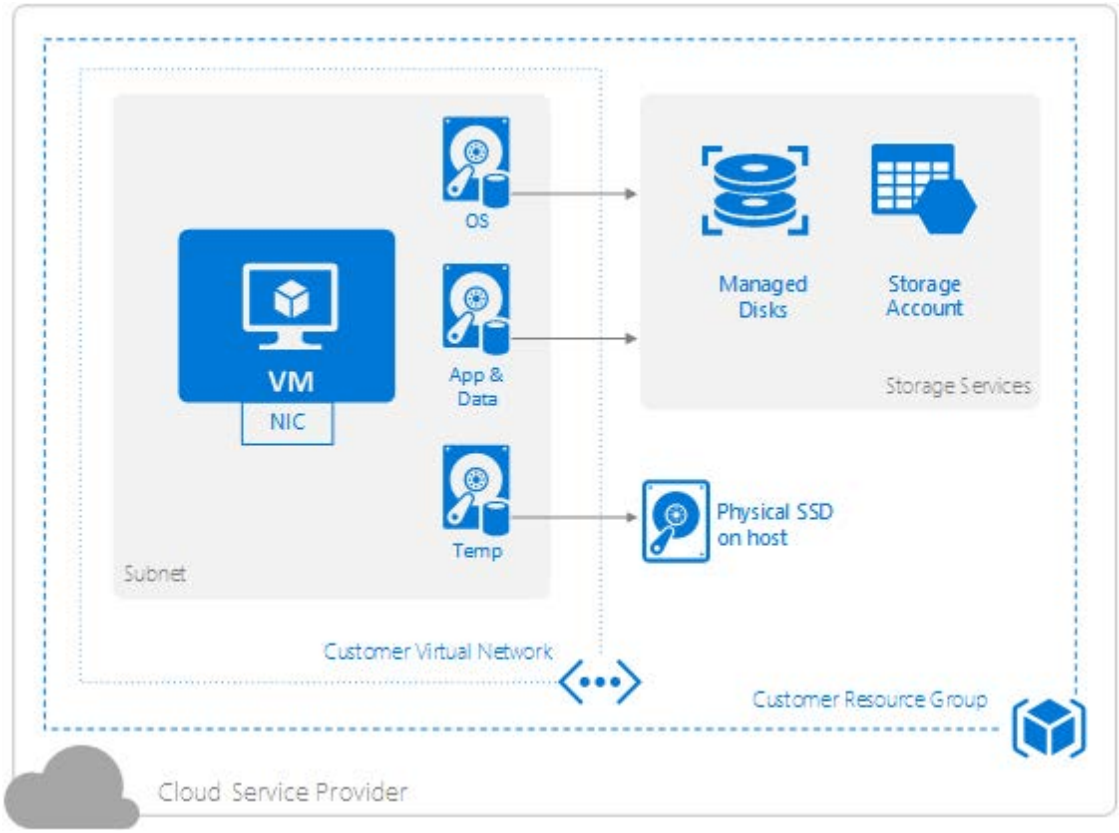

8 Deployment Architecture References

The following are a few **Azure reference architecture** examples of typical scenarios that serve as guidance for designing customer specific deployments. victor would typically live in the Web/Business/Data tiers which can be simplified to a single Application tier subnet, or Data tier subnet maintained separately based on the remote SQL configuration.

1. **VPN** - This reference architecture shows how to extend an on-premises network to Azure, using a site-to-site virtual private network (VPN). Traffic flows between the on-premises network and an Azure Virtual Network (VNet) through an IPSec VPN tunnel.

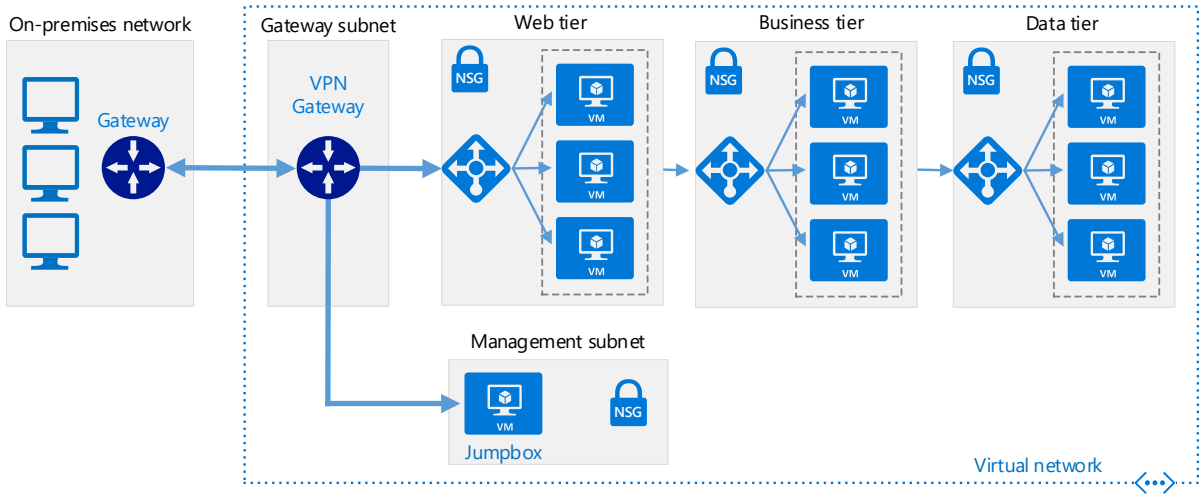

2. **DMZ Private** - This reference architecture shows a secure hybrid network that extends an on-premises network to Azure. The architecture implements a DMZ between the on-premises network and an Azure virtual network.

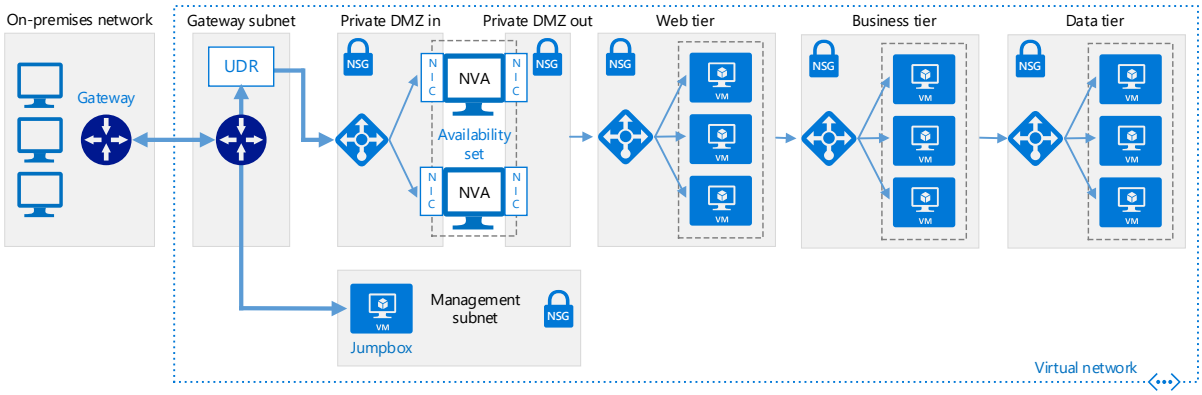

3. **Windows N-Tier** - This reference architecture shows how to deploy VMs and a virtual network configured for an N-tier application with SQL Server redundancy.

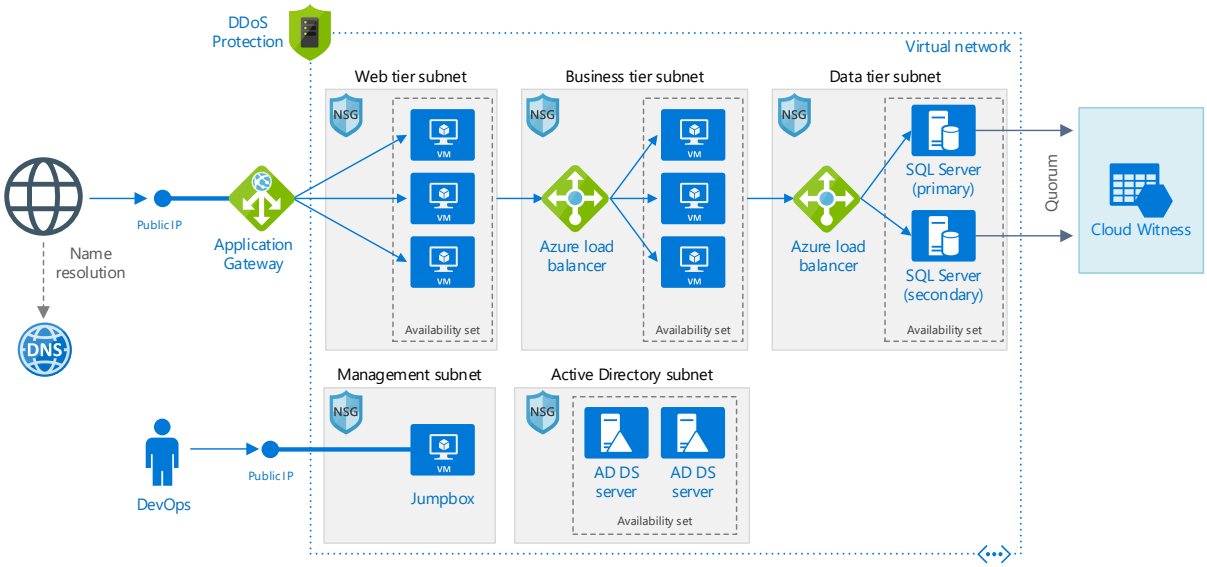

4. **Windows Multi Region** - This reference architecture shows running an N-tier application with SQL Server Always On in multiple Azure regions, in order to achieve availability and a robust disaster recovery infrastructure.

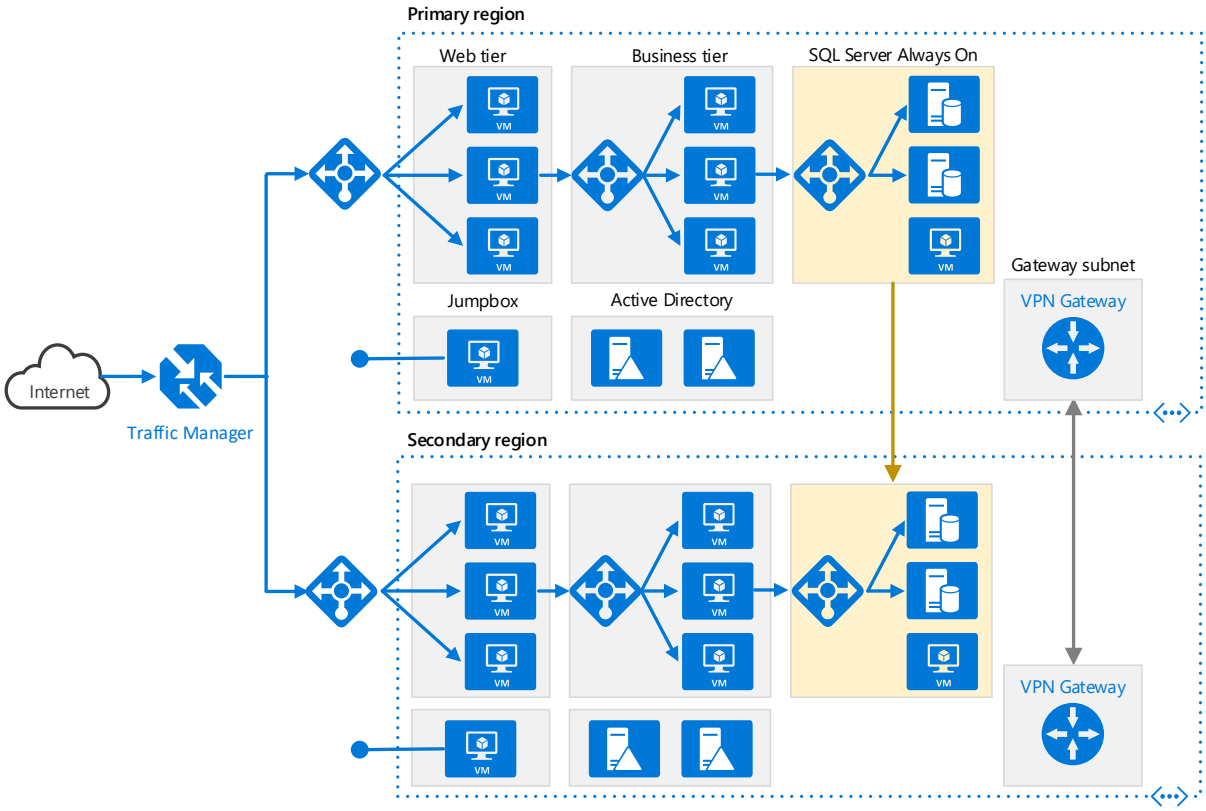

9 victor Application Licensing

victor provides a hardware flexibility license option for cloud infrastructure support that does not bind software licensing to physical or virtualized hardware instances. victor application instances can be individually licensed and deployed to run on cloud infrastructure without license validation service interruption.

victor can still be installed and traditionally licensed on a cloud VM instance without a hardware flexibility license, with the limitation that the license validation will rely on the UUID lifecycle of the VM instance.

For more information about victor hardware flexibility licensing, consult your local or regional American Dynamics sales team representative.

10 victor VMS Network Security

For purposes of network security planning, all ports and protocols used by the victor Client are detailed in the latest published version of “victor-videoedge-port-structure_rc0_kb_lt_en.pdf”, available via **www.americandynamics.net** with partner login access.

Create security group policy rules within the isolated and protected network, for restricting ingress/egress traffic to only servicing authorized users and end-point client connections required for victor.

- **Azure Network Security Groups**
- **AWS Security Groups**
- **GCP Firewall Rules**

11 victor Cloud Deployment Recommendations

Think and plan before deployment, the Cloud will not fix on-premise system problems.

1. On-Premise victor customers planning to migrate to cloud.
    a. Perform full system audit to assess health and performance of existing system.
    b. Inventory entire estate to identify all remote client and endpoint devices including integrations for cloud readiness connectivity.
    c. Start database clean-up, identify host vs controller-based events for reprograming, and archive whatever is not required to move to the cloud.
    d. If victor server is using local static IP, move to DHCP and establish DNS hostname alias, re-network remote client and endpoint devices accordingly.
    e. Upgrade all required software and firmware first on-premise, harden security and stabilize system before moving to cloud.
    f. Stage and test cloud environment to receive live cutover migration plan.

2. victor hostname DNS alias must be in place prior to starting victor application server installation on the VM instance. Approach deployment as if setting up a redundant system to establish proper communication model for on-premise remote client and endpoint devices to connect to the victor cloud VM instance. This will facilitate even a single server deployment to be moved to another zone or region and position for HA/DR readiness.

3. Provisioning of the VM instance for victor application server installation as per CSP recommended practices for typical Microsoft server application deployments.
    a. Windows Server 2016 on boot SSD disk, mounts as C drive.
    b. SQL Server 2016 on data SSD disk, mounted as 2nd drive
    c. victor on data SSD disk, mounted as 2nd drive
    d. Move Windows paging and TempDB to local temp SSD disk, mounts as D drive. This offloads a significant number of IOP from your OS & Data SSD persistent disks.

12 Cloud Training & Support

12.1 Cloud Platform Training & Certifications

Each cloud service provider offers a broad range of role-based training and certifications for building expertise with cloud skills and best practices for their respective platforms.

- **Azure certifications**
- **AWS certifications**
- **Google Cloud certifications**

12.2 victor Application Support

American Dynamics provides pre and post sales technical support for victor, as well Professional Services to assist customers with their cloud deployment strategies and system migrations.

- Pre-Sales support: https://www.americandynamics.net/Support/Contact_Pre-Sales_Support.aspx
- Technical Support: https://www.americandynamics.net/Support/contact_technical_support.aspx
- Professional Services: Can be arranged through your reginal area sales manager:
    - N America: https://s3.amazonaws.com/ad-downloads/pricelists/AD_Contact_List.pdf
    - Latin America: https://s3.amazonaws.com/ad-downloads/pricelists/AD_Latam_Brazil-Sales_contact.pdf
    - EMEA/APAC: https://www.tycosecurityproducts.com/ProfessionalService.aspx
      Email: ps.emea.apac@jci.com

13 Cloud Support FAQ

Q: Does victor IaaS support Micro-Services and application as a service capability?
A: Not today, but this is in our roadmap for the victor releases targeted for launch in 2020.

Q: Will moving to the cloud for victor require a full replacement / migration of my existing victor?
A: No, the roadmap for cloud support will shift the existing victor architecture to the new cloud platform over the next 12 months of our product launch and updates. A hybrid of the existing platform and new will be supported during this time where the clients won't know the difference between the platforms to help manage this transition.

Q: What is CrossFire?
A: CrossFire is the native server platform behind victor that is based on the Microsoft .NET Framework. This platform will be shifting to CrossFire+ and leverage .NET Standard and Core from Microsoft which will support Containers, Dockers and true Micro-Services that leverage all the advantages of cloud while additionally removing the dependency on the Microsoft Windows server OS.