Deployment Guide:

Deploying OnGuard® Systems in an IaaS Cloud Environment

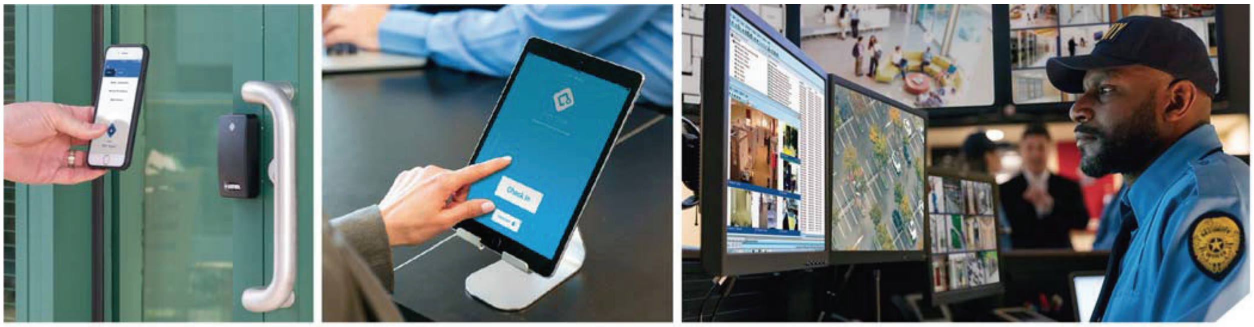

Deploying OnGuard® Systems in an IaaS Cloud Environment

Contents

Table of Illustrations

OnGuard® Deployment Guide:
Deploying OnGuard in an IaaS Cloud Environment

1 Overview

This document is a high-level design and reference guide for deploying the LenelS2 OnGuard Security Solution within a cloud infrastructure environment. Attention is given to general principles applicable to any cloud infrastructure. Specific detail is provided on two of the most popular public cloud environments, Microsoft® Azure® and Amazon® Web Services® (AWS®).

As a reference architecture, this document will describe specific deployment configurations that has been tested by LenelS2 and are considered best practices. Target audiences include LenelS2 resellers and OnGuard system managers and administrators. End users are strongly encouraged to consult with their reseller partners before undertaking a deployment of OnGuard in a cloud environment.

2 The Cloud Environment

Cloud computing environments have revolutionized the IT industry over the last decade. Cloud computing enables convenient, on-demand access to computing resources, reducing or removing the need for organizations to invest in costly server infrastructure and the personnel to operate them.

Because there is no need to purchase the computing equipment, cloud resources are delivered and purchased using a services model. Users consume and pay for computing, networking, and storage on an ongoing fractional basis, for only what they need, when they need it.

The delivery of computing services can take a number of different forms. For our purposes, three are relevant: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS).

IaaS Model

In an IaaS model, a cloud provider hosts the infrastructure components including servers, storage and networking hardware and software. An end user leases these components, typically on a Virtual Machine (VM) basis, and then installs and runs his or her applications within this environment. IaaS provides the tools and services necessary to assist the end user in maintaining and managing infrastructure, including the ability to adjust the amount and type of the component resources used to best match the applications' needs.

PaaS Model

PaaS extends the IaaS model by offering end users access to supporting software components, such as middleware and application runtime environments. This reduces the amount of software that the end users would otherwise need to develop or provide in order to run their application suites.

SaaS

With SaaS, host providers extend the concept an additional step by hosting and managing the applications themselves, and offering them to end users as a service which can be leased in the same manner as the infrastructure. SaaS users do not need to install, license, or maintain any applications themselves. The host provider is responsible for all maintenance and version management. End users simply pay a recurring fee for access to and use of the software that they need, when they need it.

Cloud Service Providers by Model

There are many cloud providers for both IaaS and PaaS solutions, with data centers spanning the globe. SaaS providers very often are application-specific vendors and the choice of cloud provider is determined by the application vendor.

Cloud service providers can offer advanced capabilities for security and performance, combined with network services such as load balancing and WAN optimization services. Beyond the choice of cloud computing model, consideration should be given to a provider's overall capabilities. Specific evaluation criteria may include assessing a provider's number and location of data centers, their ability to support various WAN traffic architectures and bandwidths, their data security and compliance models, and their support of legacy application environments.

3 The OnGuard® Platform & The Cloud

OnGuard is a robust, globally scalable security platform with extensive integrations to other platforms and applications. It is a time-tested application suite with an extensive feature set and an open API for creating multi-partner solutions. It is a client/server-based solution with global distribution properties. The following diagram illustrates the extensive scope and scalability of an Enterprise-class OnGuard installation.

Figure 1. OnGuard Capabilities

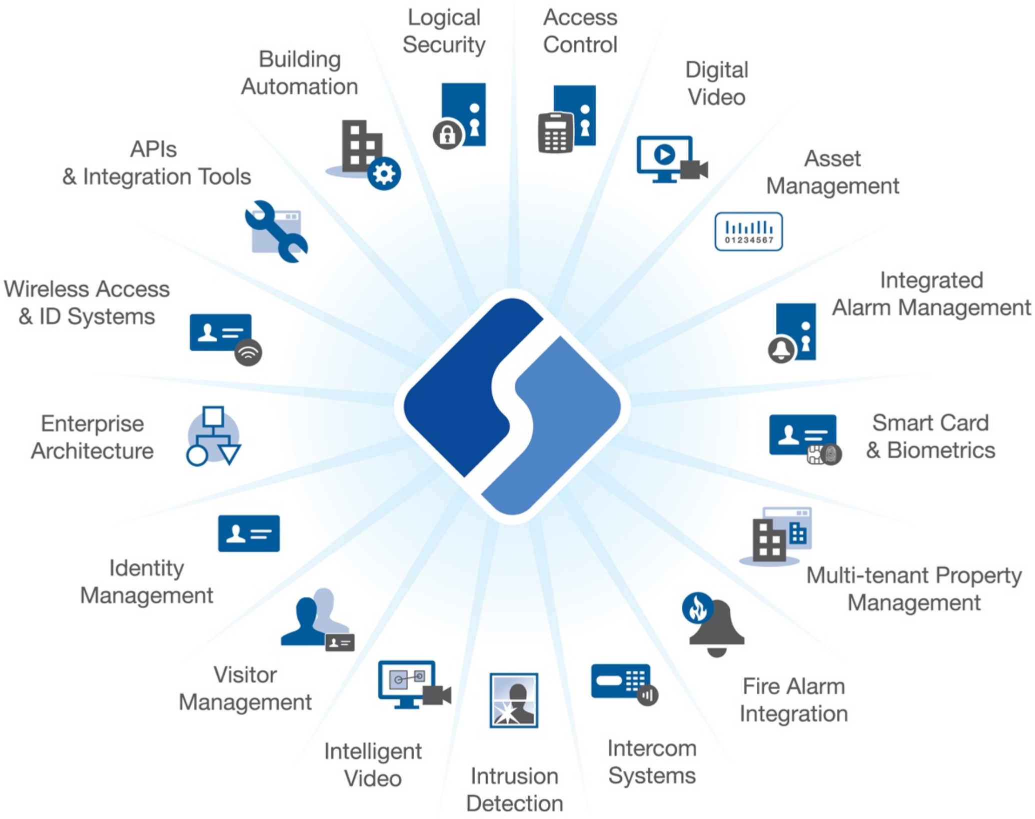

While OnGuard® is architected as a client-server solution, it can certainly be made to work in a cloud environment and can take advantage of many of the capabilities that make the cloud a compelling alternative to on-premise data centers.

Deploying OnGuard in a cloud environment can lower capital costs through the efficient lease of virtualized resources in place of the purchase of dedicated hardware. Operating costs can be lowered by leveraging cloud data center personnel to manage the virtualized and shared hardware infrastructure. Highly distributed and global deployments can be simplified by taking advantage of the fact that large cloud providers have already invested in global footprint data centers. Cloud environments also offer high levels of built-in resiliency and redundancy. They can simplify IT compliance requirements by offering such compliance as a core competency of the provider.

With these benefits in mind, careful attention should be paid to the tradeoffs and potential risks that the use of cloud hosting can impose. Principal among these is the reliability and overall throughput of the network connections between the premise(s) to be secured and the cloud infrastructure. Panels, readers, cameras, and other security sensors and devices must necessarily remain on the secured site. And while these devices, when properly deployed, can function autonomously in the event of a Wide Area Network (WAN) loss, the ability to properly monitor and manage them can be seriously curtailed. Additionally, traditional, full featured 'thick' client applications may not perform acceptably across the higher latency, lower bandwidth links that often accompany onsite client-to-cloud server connections. Alternative approaches may be required, including the use of internet-friendly web applications and/or remote desktop applications. Large internet and cloud hosting providers often provide high speed, low latency options for cloud connectivity but these can be expensive. Appropriate performance profiling should accompany any decision to deploy your security solution in a cloud environment.

In the simplest sense, deploying OnGuard® into a cloud environment is very similar to an on-premise deployment using Virtual Machines (VMs). OnGuard is well suited to leverage the IaaS and selected PaaS capabilities of large cloud providers as long as its client-server roots are kept in mind when designing such a deployment. The following sections outline the best practices for creating such a deployment. These best practices should apply equally to any cloud environment, whether public or private in nature. Some specific consideration is also given to two of the most popular public cloud providers, Amazon AWS and Microsoft Azure.

4 OnGuard using Infrastructure as a Service

Many of the considerations for running OnGuard in an IaaS cloud environment are similar to those that must be addressed when deploying the same system in an on-premise environment, especially if that environment is highly virtualized. Attention must be paid to the capacity of the compute (server) instances, the connectivity between the server instances, between the servers and clients, and between the servers and on-premise equipment. The choices made will be strongly influenced by the size of the OnGuard system being considered. Three different topologies are described in this document, one each for the OnGuard tiers 32 ES/ADV, Pro, and Enterprise.

Independent of the OnGuard tier being deployed, there are some general guidelines that can be applied.

Utilize a strong Virtual Machine (VM) for the OnGuard Application Server. Many cloud providers offer different processor types and memory configurations for their Virtual Machines. These configurations can be specialized for bursty workloads, computationally intensive tasks, high throughput, etc. For OnGuard, a general purpose processor class is recommended. The mimimum hardware requirements found in the OnGuard Specifications Sheets provide a good starting point.

For small to medium installations, combine the Application, Database and Comm Servers in one VM. This approach is a common practice in on-premise OnGuard installations and is certainly acceptable in a cloud setting where virtual machines are engineered to be highly available, data redundant and resilient to outages.

When using separate VMs for the Application, Database, and Comm Servers, consider colocation. Low latency connections between these services are critical to assuring a high performing OnGuard® installation. (See the section on *Performance Profiling* for acceptable system latencies.) Placing these services in different VMs may help with resilience and performance tuning so long as the connectivity between them is robust. The VMs supporting these services should be in the same data center and the same virtual subnet.

Create a Virtual Private Cloud for secure site-to-cloud connectivity. Unless there is a specific requirement for the OnGuard system to be accessible over the public Internet, on-site clients and security equipment should communicate with the cloud-based OnGuard servers via a secure connection. The best practice is to use on-premise VPN hardware, connected to a virtual VPN/Cloud gateway from the Cloud provider. Your organization's on-site network addressing schema should be extended across the VPN connection and be used by the virtual subnet(s) hosting the OnGuard VM(s). This Virtual Private Cloud approach simplifies the overall design and administration of your IaaS OnGuard system.

Use OnGuard Browser-based Clients wherever and whenever possible. The new OnGuard browser-based clients are built for Internet connectivity and are much less sensitive to network delays than the older, Windows® based thick clients. If, for whatever reason, you must use one or more thick clients then there are two choices for assuring performance. One is to run the thick client in the Cloud in a VM and use a remote desktop console on-site, communicating with the cloud client via a Remote Desktop Protocol. The other approach is to use a high performance WAN link. The latency target should be in the 50ms range. Amazon AWS provides this option via their Direct Connect® service. Microsoft Azure has a similar offering called ExpressRoute®.

Follow existing OnGuard practices to scale the system. OnGuard is designed as a 'scale up' architecture meaning the common Web practices of adding servers and load balancers to improve performance are not applicable. Scaling OnGuard in the cloud can be achieved using the same mechanisms that have been developed for on-site deployment. OnGuard servers can be scaled *up* by selecting more powerful processor classes and VMs. The good news is that it is easy to experiment with this approach in a cloud setting as no hardware purchases are required to perform such tests. There is one area where OnGuard can be 'scaled out'. Scale *out* the capacity of the on-site security equipment, such as the number of panels supported, by adding additional Comm Servers in one or more additional Virtual Machines. Just remember to consider colocation of these additional Comm Servers with their respective Application and Database Servers. For geographic scale, leverage the OnGuard Enterprise Edition® Regional database server solution.

4.1 OnGuard® IaaS Reference Designs

This section provides design archetypes for each of the following 'system size' versions of OnGuard:

- OnGuard 32ES/ADV system
- OnGuard Pro system
- OnGuard Ent system

Using the principles of the previous section, each design builds on the preceding in terms of the complexity and the potential scale of the deployment.

4.1.1 OnGuard 32ES/ADV Architecture for the Cloud

For small OnGuard system installations of up to 256 readers and 10 clients, a single VM installation, containing all of the OnGuard services, is the most efficient and cost effective approach. The following diagram documents this approach.

Figure 2. OnGuard 32ES/ADV Architecture – Thin Clients

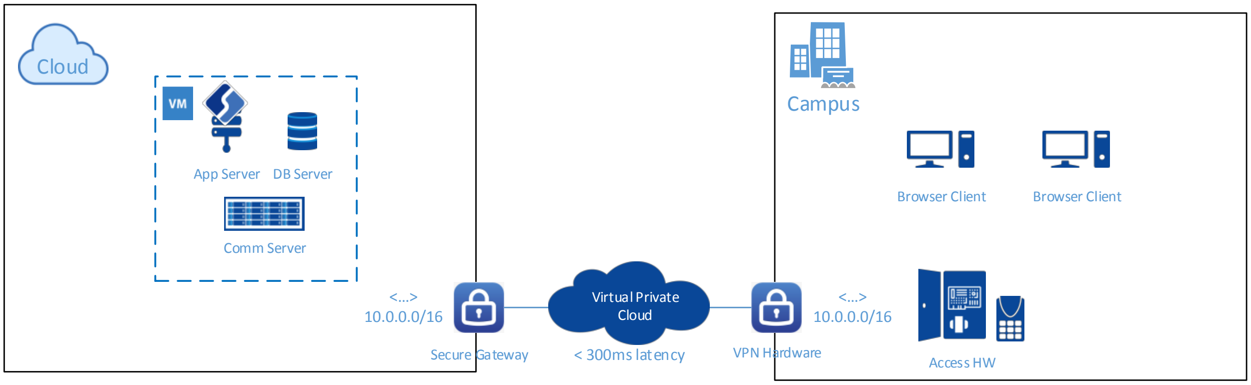

In this scenario, a Virtual Private Cloud is created that extends the campus network to the cloud via secure VPN connections. In the cloud environment, a single VM is instantiated that hosts all of the OnGuard server components. Most cloud providers design redundancy into the physical hosting of their VMs, so this configuration in the cloud is more resilient than a single physical server operating on the campus.

The following attributes and considerations apply to this scenario:

OnGuard Version: OnGuard system version 7.4 or later is supported.

Virtual Machine: A general performance processor such as an Intel® Xeon® E5 (“Ivy Bridge” or “Sandy-Bridge”) with 8GB of RAM is recommended. Microsoft® Windows 10 Embedded IoT® is recommended. For other compatible versions of Windows, see the *Operating System Compatibility Chart* on the Lenel Partner Center website.

Cloud to Site Connectivity: Most cloud providers offer a Secure Gateway/Virtual VPN service to secure the connection on the cloud side of the WAN connection. On the user premise a hardware VPN device is recommended. Bandwidth requirements are difficult to forecast, but network latency tolerances can be bounded. For this scenario, network latencies should be below 300ms in order to assure reasonable client response and acceptable performance for the communication between the Comm Server and the on-site access control panels.

On-Site Access Control Equipment: There are no cloud-specific restrictions regarding the supporting access control hardware. See the *Lenel Access Control Hardware Compatibility Chart* on the LenelS2 Partner Center portal.

Client Applications: In the simplest deployment scenario it is envisioned that all on-site access to the system be performed using the OnGuard® web clients. This approach provides the best user experience and mitigates the need to invest in more costly site-to-cloud network capacity. If thick clients are required, then the most cost effective option is often to run the thick clients in the cloud environment, in a client suitable VM, and create an RDP connection from the cloud to an RDP console on-site. Depending on the number of open applications, and the degree of graphics processing being performed on the thick client, the cloud client VM can be sized from an Intel I-3 with 4GB RAM to an Intel E5 with 8GB RAM.

Figure 3. OnGuard 32ES/ADV Architecture – Thin & Thick Clients

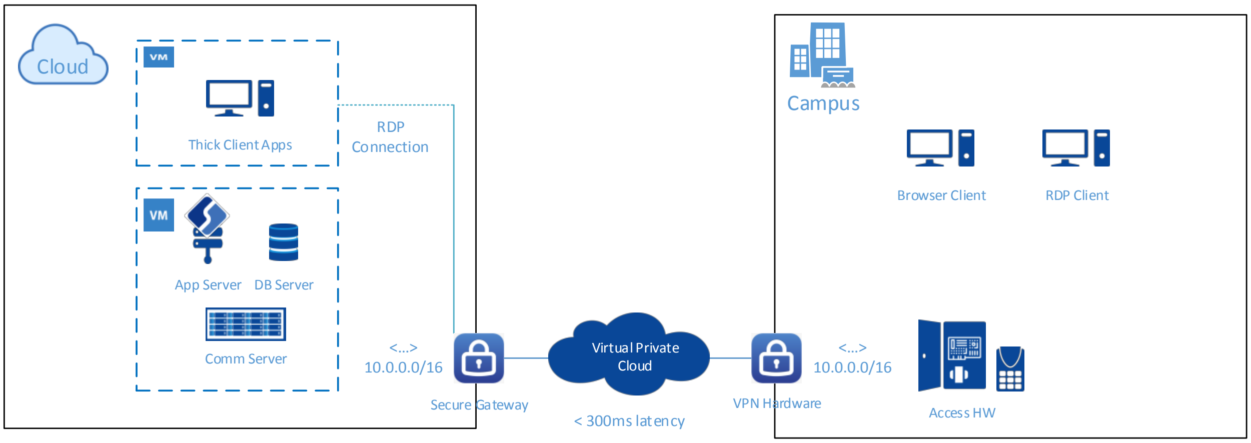

4.1.2 OnGuard Pro Architecture for the Cloud

For larger OnGuard system installations that exceed 256 readers and 10 clients but that are not highly geographically distributed, the OnGuard server components can be divided into multiple VMs in order to optimize performance, all contained within a single cloud datacenter or geographical region. The degree to which multiple VMs are necessary is a function of the size and degree of traffic flow within the system and should be determined empirically.

Figure 4. OnGuard Pro Architecture

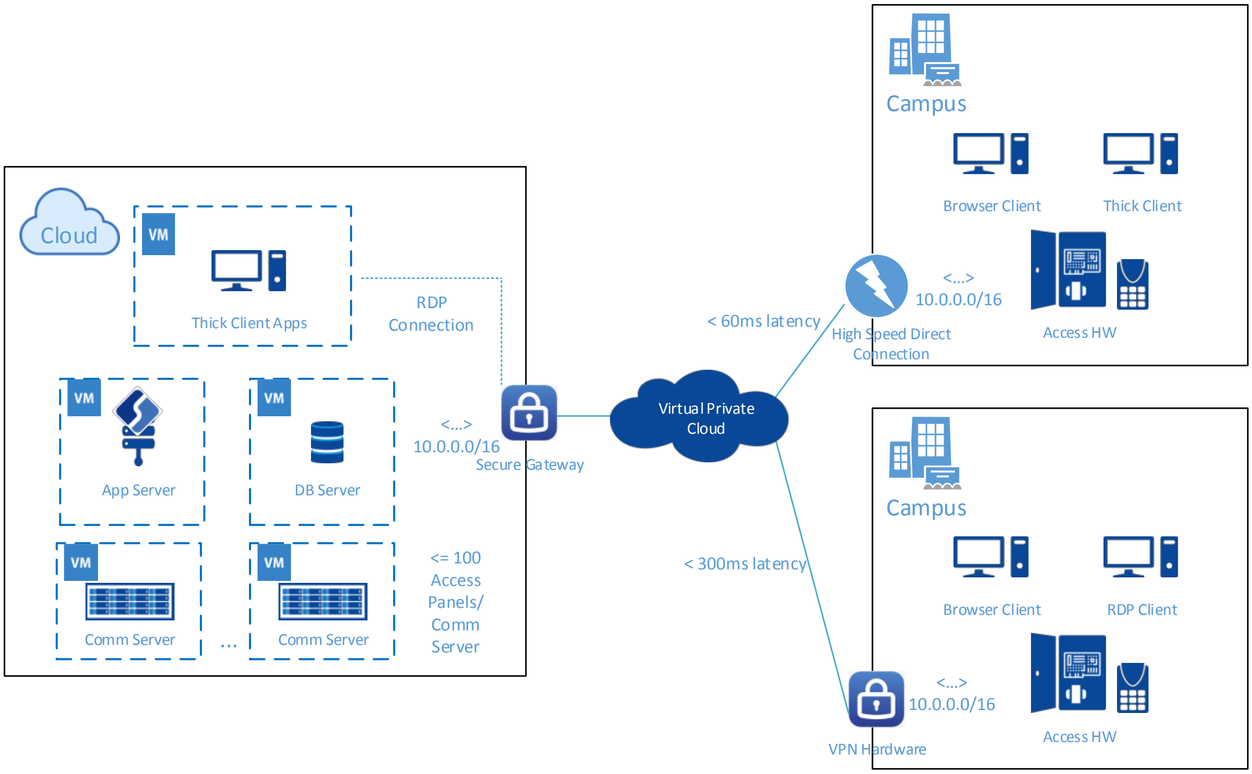

In this scenario, a Virtual Private Cloud is created that extends multiple campus networks to the cloud via secure VPN connections and/or dedicated high speed connections. In the cloud environment, the OnGuard® Application Server, Database Server, and Comm Servers are instantiated on separate VMs with the same datacenter/geographic region. The specific VMs chosen can be tuned to the specific needs of the server and the types of traffic that it processes. Whether each server needs to be placed in its own VM should be determined through performance testing and availability profiling.

The following attributes and considerations apply to this scenario:

OnGuard Version: OnGuard system version 7.4 or later is supported.

Virtual Machines: General performance processors such as those in the the Xeon E5 class (“Ivy Bridge” or “Sandy-Bridge”) or later generation are acceptable but higher cache and RAM values should be considered (up to 15MB and 16GB respectively). Comm Servers should be initially sized based on access control panel counts. 100 access panels per Comm Server is a good rule of thumb, keeping in mind that traffic volumes vary from panel group to panel group. Microsoft® Windows Server® 2012 or 2016 should be installed. For other compatible versions of Windows, see the *Operating System Compatibility Chart* on the LenelS2 Partner Center portal.

VM to VM Connectivity: When placing the OnGuard® server components into separate VMs, special attention needs to be paid to the VM to VM connectivity and virtual network latency. Application Server to DB Server, Application Server to Comm Server, and Comm Server to DB Server should maintain connection latencies of between 5ms and 40ms, with the low end of that range highly preferred. This implies that these server VMs should be colocated within the same cloud datacenter, possibly within the same physical rack space if the provider supports such an option.

Cloud to Site Connectivity: Most cloud providers offer a Secure Gateway/Virtual VPN service to secure the connection on the cloud side of the WAN connection. On the user premise a hardware VPN device is recommended. Bandwidth requirements are difficult to forecast, but network latency tolerances can be bounded. For Comm Servers, thin clients running OnGuard web clients, and remote desktop clients network latencies should be below 300ms in order to assure reasonable client response and acceptable performance for the communication with the on-site access control panels. If thick clients must be deployed on-site, then higher speed network connections should be employed that assure latencies at or below 60ms. Examples of higher speed links include ExpressRoute (Microsoft Azure) and Direct Connect (Amazon AWS).

On-Site Access Control Equipment: There are no cloud-specific restrictions regarding the supporting access control hardware. See the *Lenel Access Control Hardware Compatibility Chart* on the LenelS2 Partner Center portal.

Client Applications: In the scenario pictured, a mix of thin, remote desktop, and thick clients are deployed. Where thick clients are considered necessary, then high speed site-to-cloud links should be employed as discussed under the Connectivity subsection, above. For cloud-based thick clients utilizing RDP connections, the cloud client VM can be sized from an Intel I-3 with 4GB RAM to an Intel E5 with 8GB RAM depending on the number of running applications and the level of graphics processing being performed.

4.1.3 OnGuard Enterprise Architecture for the Cloud

For larger OnGuard system installations that are geographically distributed, the OnGuard Enterprise system version offers a Master/Region database server model that is a good fit to cloud providers supporting a multi-datacenter/multi-regional network. The following diagram documents this approach.

Figure 5. OnGuard Enterprise Architecture

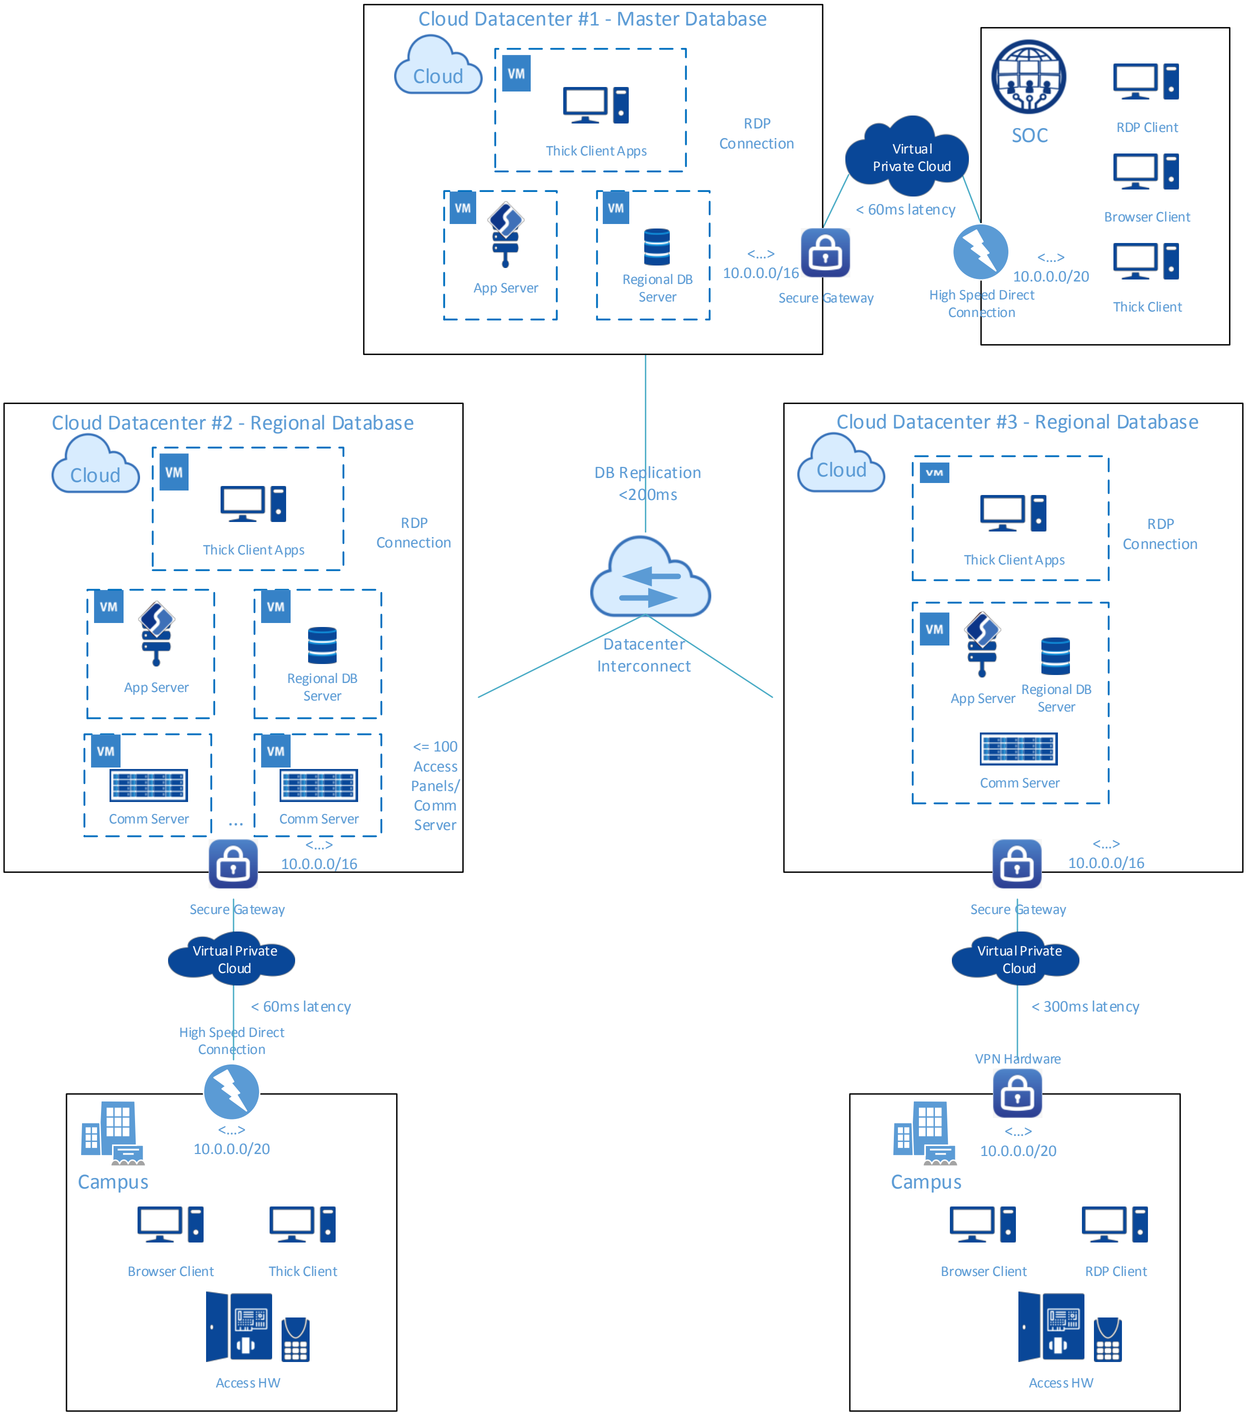

In this scenario, a Virtual Private Cloud is created that extends multiple campus networks to multiple cloud datacenters via secure VPN connections and/or dedicated high speed connections. In the cloud environment, the OnGuard® server components may be placed in separate VMs or combined, based on size and level of activity at the sites to which they connect. One of the cloud datacenter sites is setup with the OnGuard® Master Database. The other sites are configured as Regional Databases. A coherent view of the entire set of secured campuses is achieved through replication of the Regional Databases to the Master.

The following attributes and considerations apply to this scenario:

OnGuard Version: OnGuard system version 7.4 or later is supported.

Virtual Machines: General performance processors such as those in the Xeon E5 class ("Ivy Bridge" or "Sandy-Bridge") or later generation are acceptable but higher cache and RAM values should be considered (up to 15MB and 16GB respectively). Comm Servers should be initially sized based on access control panel counts. 100 access panels, per Comm Server is a good rule of thumb, keeping in mind that traffic volumes vary from panel group to panel group. Microsoft Windows Server 2012 or 2016 should be installed. For other compatible versions of Windows, refer to the the *OnGuard Operating System Compatibility Chart* on the LenelS2 Partner Center portal.

VM to VM Connectivity: When placing the OnGuard server components into separate VMs, special attention needs to be paid to the VM to VM connectivity and virtual network latency. Application Server to DB Server, Application Server to Comm Server, and Comm Server to DB Server should maintain connection latencies of between 5ms and 40ms, with the low end of that range highly preferred. This implies that these server VMs should be colocated within the same cloud datacenter, possibly within the same physical rack space if the provider supports such an option.

Cloud Datacenter-to-Datacenter Connectivity: Depending on the location of the campus environments to be secured, the cloud datacenters to which they attached may be widely distributed geographically in order to assure the performance of other aspects of the system. To assure suitable performance for the database replication function, the datacenter-to-datacenter connectivity should have a latency of less than 200ms.

Cloud to Site Connectivity: Most cloud providers offer a Secure Gateway/Virtual VPN service to secure the connection on the cloud side of the WAN connection. On the user premise, a hardware VPN device is recommended. Bandwidth requirements are difficult to forecast, but network latency tolerances can be bounded. For Comm Servers, thin clients running OnGuard browser-based web clients and remote desktop clients network latencies should be below 300ms in order to assure reasonable client response and acceptable performance for the communication with on-site access control panels. If thick clients must be deployed on-site, then higher speed network connections should be employed that assure latencies below 60ms. Examples of higher speed links include ExpressRoute® (Microsoft Azure®) and Direct Connect® (Amazon AWS®).

On-Site Access Control Equipment: There are no cloud-specific restrictions regarding the supporting access control hardware. See the *Lenel Access Control Hardware Compatibility Chart* on the LenelS2 Partner Center portal.

Client Applications: In the scenario pictured above, a mix of thin, remote desktop, and thick clients are deployed. Where thick clients are considered necessary, then high speed site-to-cloud links should be employed as discussed under the Connectivity subsection, above. For cloud-based thick clients utilizing RDP connections, the cloud client VM can be sized from an Intel I-3 with 4GB RAM to an Intel E5 with 8GB RAM depending on the number of running applications and the level of graphics processing being performed.

5 Additional Deployment Considerations

5.1 On Premise Hardware

Access control panels, readers, and related security equipment, must, by definition, remain on-site. All of the server components, however, should reside in the cloud. It is certainly possible to create hybridized architectures of cloud-based and on-premise-based infrastructure, but care should be taken to assure appropriate groupings of server types and network connectivity, as indicated in the preceding diagrams.

Comm Servers deserve special mention as they are the most common scale-out component of an OnGuard® system and might seem a candidate for on-premise placement. **The critical design point for a Comm Server is not its connectivity to the access control panels but rather to the Application and Database Servers. Colocation of these components is strongly encouraged for a high performing installation.**

Client workstations should leverage the new OnGuard browser-based clients whenever possible. These applications are designed for internet connectivity and are far more tolerant of varying network latencies. As described in the preceding sections, if thick client apps are considered necessary, there are two options to consider. The first is to run those client applications in the cloud and use remote desktop protocols and applications to display them on-premise. The second is to assure the necessary network bandwidth and quality of service attributes through the use of higher quality internet connections. Many cloud providers offer dedicated cloud connections for these types of situations.

On-premise badge printing can be problematic with cloud-hosted systems. OnGuard system v7.5 provides a service that can be installed locally that allows local printers to print badges from the cloud.

5.2 OnGuard® Database Server

The OnGuard Database is the critical resource of the system. Every OnGuard service should be “close” to the database, meaning network latencies should be minimized through the use of an appropriately sized virtual machine for the database itself. Scaling up or down the database VM is a major performance variable. Guidance is provided in earlier sections regarding the sizing of this VM. Cloud environments make experimenting with different VM types a simple and fast proposition and it is worthwhile to do so during the deployment phase.

Microsoft® SQL Server® and SQL Server Express® are supported for cloud-based deployments of OnGuard. A full list of compatible versions is available on the Lenel Partner Center portal. Oracle® Database Server has not been tested for cloud deployment and is not supported at this time.

With the release of OnGuard system v7.5, cloud deployments in the Microsoft Azure environment now support the Azure SQL® database. For some organizations this service-based offering may represent an attractive alternative to purchasing a full Microsoft SQL Server license. Refer to the Microsoft Azure home page for additional information and pricing options.

5.3 Network

A site-to-site connection (S2S) joins a cloud virtual network to an on-premise network via an IPSEC tunnel established through the use of a (preferably) hardware-based VPN device. Organizations can domain-join VMs in the cloud network and leverage their on-premises domain name servers (DNS) for address resolution between on-premise and cloud assets. The persistence, cost, and security profile of a S2S solution make it the most appropriate option for an OnGuard system IaaS implementation.

Configuring an S2S connection requires the following high-level components:

1. A cloud-based virtual network, whose address space must not overlap with the private addresses of the on-premise network
2. A cloud-provided virtual network gateway with an external IP address used for connections to the on-premises network
3. A hardware VPN device acting as a local (on-premise) network gateway, configured with an external IP address
4. A site-to-site Connection that joins the Virtual Network Gateway and the Local Network Gateway via a trusted, shared key.

ExpressRoute and AWS Direct Connect are two offerings, from Microsoft and Amazon respectively,which offer a direct connection to the cloud, bypassing the public internet. The reliability and speed of such connections make these services a good fit for scenarios involving thick client applications, data migration and disaster recovery.

5.4 Video Considerations

Managing enterprise video in a cloud environment can be a complicated task. Video rates and resolutions, retention periods for storage, and the degree to which video will be constantly monitored or only accessed upon event, all have to be carefully considered against the cost models of the various cloud providers under consideration. As an example, in Azure and AWS environments, storage of video is reasonably inexpensive, but transmitting the video from cloud storage to on-premise clients can carry considerable cost. The OnGuard® system does not currently have a native solution for storing video in the cloud. These considerations, combined with the sizeable investment many customers have made in on-premise video systems, suggests that one approach is an architecture that uses an edge-based storage design along with a hybrid approach to video viewing and retrieval.

The following diagram illustrates this concept:

Figure 6. Edge-Based Video Design with Multiple Viewing Options

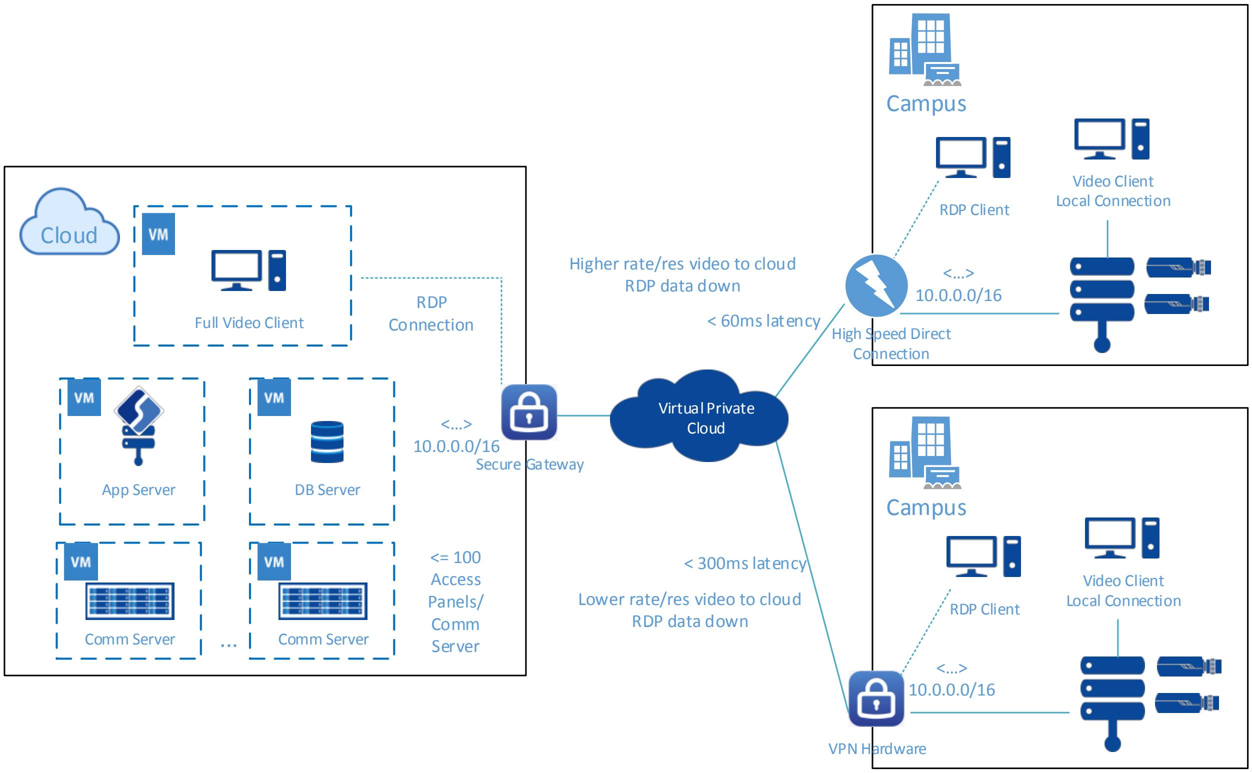

This scenario follows the same Virtual Private Cloud model as the previous ones, but unlike the access control server components, the video server/recorder units remain on-premise. The video server/recorder units may be accessed and managed via the cloud infrastructure but the video content itself is not stored in the cloud.

Video monitoring and retrieval may be performed in one of several ways. The video may be retrieved directly from the recorder when the client is within the same campus or LAN environment. A similar approach may be used across campuses when those campuses share network connections independent of the cloud connection and the network speed is sufficient to support such retrieval. When the campus interconnect is supported only via the Virtual Private Cloud connection, the best approach may be to run the video client in a VM within the cloud infrastructure and use an RDP connected client on the receiving campus to view the video. This type of approach makes the best sense when, as mentioned above, the cloud provider charges for the transmission of the video from cloud to campus (but not for receiving same). This approach also makes sense when bandwidth is limited or there are multiple monitoring clients on a campus since only the video display data must be transmitted and not the full video stream.

The following attributes and considerations apply to this scenario:

OnGuard® Version: OnGuard system version 7.4 or following is supported.

Video Servers/Recorders: In principle, any on-premise video recording solution native to OnGuard or supported via the OnGuard OAAP program should be compatible with this approach. Care should be given to the configuration of recording stream rates and resolutions to assure that the combinations chosen are compatible with the capabilities of the campus-to-cloud and campus-to-campus network connections.

Client Applications: In Figure 6, a mix of thin, remote desktop, and thick clients are deployed. Web browser-based or thick video clients are depicted using local and campus-to-campus connectivity to view video. These connections avoid the transmission of video to and from the cloud, minimizing the potential for additional network latency and cloud expense. It is assumed that management traffic flows through the cloud in order to set up the video connections to the edge-based servers. As a result, cloud connectivity is still required for proper traffic routing.

If only campus-to-cloud connectivity exists (i.e., there is no direct campus-to-campus network option) then it is recommended that RDP clients be deployed in the cloud and RDP clients used on the campuses for remote viewing of video from other sites. The cloud client VM should be configured with an Intel E5 or later processor and 8GB of RAM (16GB highly recommended).

5.5 Licensing Options

As of this writing, cloud-based deployments of OnGuard involve the same licensing requirements and mechanisms as on-premise systems. The license server should reside in the cloud along with all of the other server components. All cloud-based licenses require a 'file level' expiration, even if perpetual. Licensing remains End User specific.

5.6 Support Statement

Support for customers with the OnGuard® system installed within an Infrastructure as a Service (IaaS) model will be very similar to virtual environment configurations that LenelS2 supports today. Regardless of the cloud provider selected, LenelS2 support for these installations will be limited to the virtualized instance of the OnGuard Application or Component Servers and all support provided will be focused on the OnGuard system installation.

Virtualized instances of servers/workstations and networks used must meet or exceed the same recommended resource requirements as their physical counterparts in similar roles. As such, LenelS2 will support OnGuard in this environment using the same process and tools to troubleshoot and investigate problems that are part of our 'internal toolkit', as well as commercially available tools such as Wireshark.

5.7 Conclusion

Deploying an OnGuard system in an IaaS cloud environment can be an effective way of managing the capital and the operational costs of the system. In designing a cloud deployment, the scale of the system, in terms of capacity as well as geography, must be considered. Special attention should be paid to the placement of the server components of the solution. Specific guidelines to minimize variables such as connection type and network latency should be followed. LenelS2 support for IaaS installations is limited to the virtualized instance of the OnGuard Application or Component Servers and all support provided will be focused on the OnGuard system installation. Following the best practices described in this document will help ensure a successful IaaS deployment of OnGuard.

6 Reference Information

6.1 Network Latencies by Connection Type

Table 1. Network Latencies by Connection Type

Connection Type	Network Latency
Application Server to DB Server	<= 40ms; 5ms typically observed
Applications Server to Comm Server	<= 40ms; 5ms typically observed
Comm Server to DB Server	<= 40ms; 5ms typically observed
Comm Server to hardware	<= 280ms
Master DB to Regional (Replication)	<= 200ms
Thick Client to Application Server	35ms-50ms
Thick Client to DB Server	35ms-50ms

6.2 Panel Download Times

There are a number of variables involved in estimating panel download times. The following table describes two variables: network latency and credential complexity. This data should be used as guidance only and is not intended to be fully deterministic.

Table 2. Panel Download Time Estimates

Comm Server to Panel Latency	Total # Badges to be Downloaded	Max number of Access Levels per Badge	Estimated Panel Download Time
50ms	10,000	64	1:00
50ms	10,000	128	1:35
200ms	10,000	64	2:50
200ms	10,000	128	4:20
50ms	50,000	64	4:30
50ms	50,000	128	7:50
200ms	50,000	64	13:00
200ms	50,000	128	22:00

LenelS2.com

Specifications subject to change without notice.