

# **Making the Cloud Safe for Sensitive Workloads**

**Protecting Data in Motion among Data  
Centers and Servers in the IaaS Cloud**

## Introduction

The cloud provides a compelling case for cost savings, agility and operational efficiency that cannot be ignored. Executing IaaS (Infrastructure as a Service) workloads while protecting sensitive information has been challenging if not impossible in the past, but new solutions are emerging that will keep sensitive information secure in cloud environments. Customers will be able to take advantage of cloud-based servers and efficient cloud operating models while minimizing the risk of a data breach or failed audit. These solutions benefit cloud computing customers by providing the flexibility to migrate applications and servers to the environment that best fits the budget, security requirements, technical requirements and needs of the organization, without changing the underlying architecture. Flexibility is important because secure cloud computing is still in its nascent stages. Customers who adopt flexible solutions will be prepared to take advantage of new technologies and cost efficiencies in networks and cloud operating environments as they emerge.

In this paper we will explain some of the problems and new solutions that propose to make the cloud safe for sensitive workloads.


## What you will learn in this paper

- How to balance business needs, cost and security when considering cloud deployments
- Principles for optimizing and securing data centers and cloud environments
- Why the network is critical to securing sensitive information in an IaaS environment
- Recognize common problems with securing data in motion in cloud environments
- Understand the emerging technology used to solve these problems

## Principles for application and network architectures

We will start by reviewing some principles for application and network architectures. By following these principles, you can develop flexible solutions that allow changes to the solution as your requirements and costs evolve without breaking the model.

We discuss these principles by thinking about the components that make up the environment: the workload (applications, services, databases, and so on), the execution environment (physical and virtual servers, virtualization layer and management), and the network. These are illustrated in Figure 1. Each of these plays a critical role in defining your optimal solution architecture.



**Figure 1: Components of the environment**

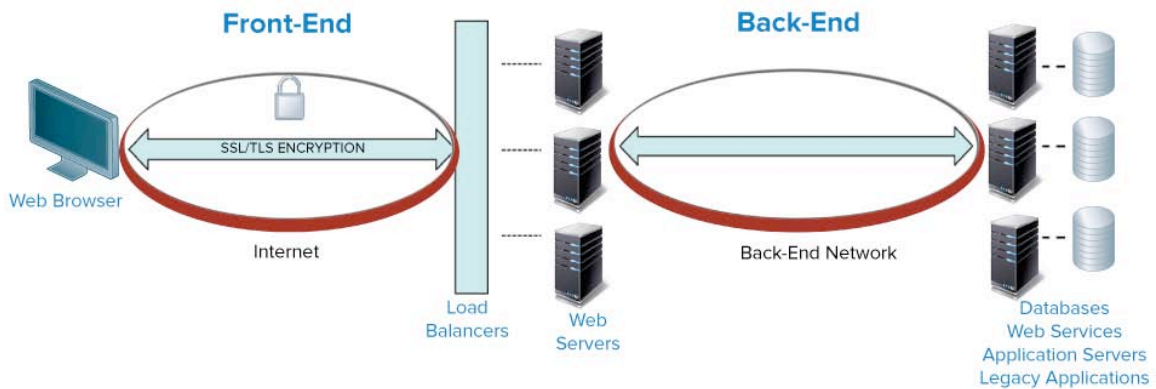
Each of the components of your workload (applications, services, databases, etc.) should run in the environment in which it is best suited to run. Legacy applications that run on specialized or older hardware should continue to run in the corporate data center. Databases and other applications that cannot be easily virtualized or that experience unacceptable performance penalties when virtualized should run on dedicated servers without being virtualized. Applications that can be virtualized should be virtualized if at all possible in order to reduce the number of physical servers and optimize server performance. Virtualization and cloud operating models should be used to optimize resource utilization and minimize operating costs. Applications with predictable and consistent resource requirements over time may run optimally in a private data center or private cloud. Virtualized applications that have variable resource requirements can run in the environment

that meets their connectivity and security requirements in the most cost-effective way.

Your execution environment should fit the needs of your applications. Ideally, the execution environment should not require a lot of changes to existing applications or networks, because changing applications to suit the execution environment increases cost and reduces agility. The ideal execution environment allows you to take advantage of cloud efficiencies, while leveraging existing data center and in-house resources to meet your business needs. Other cost factors such as capital costs, power, cooling, software licensing and maintenance costs play a role in determining a cost-optimized combination of data center, private cloud and hybrid cloud resources.

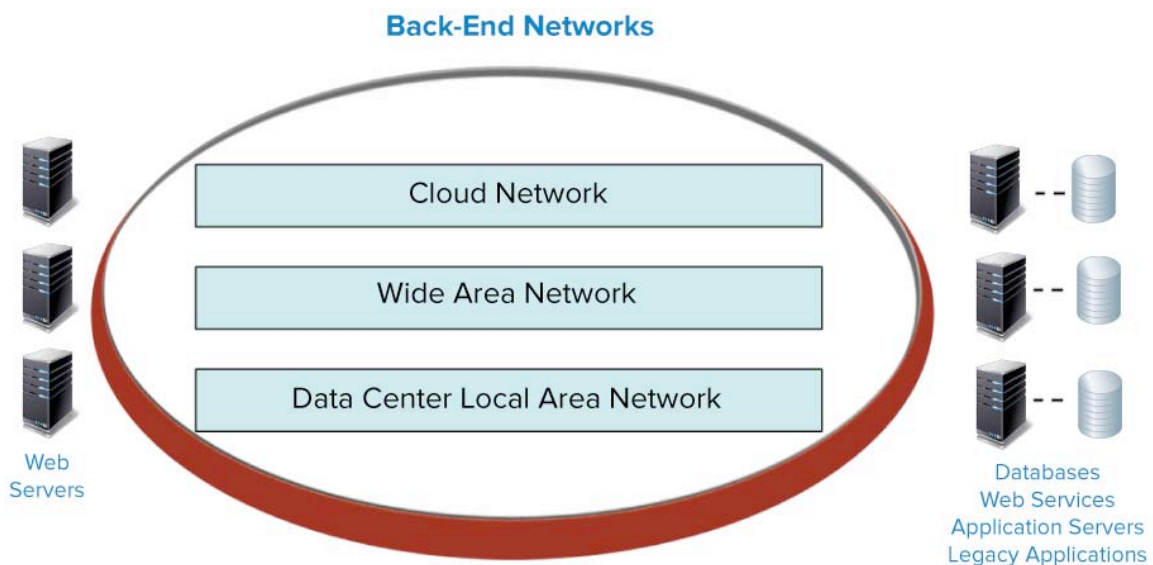
A critical component of an optimized IT environment is the network. The network connects distributed data centers, cloud environments and applications seamlessly to allow secure communication among the applications and servers in the environment. The network must be reliable and resilient to failures in order to support business-critical applications and services. The network must also provide inexpensive data transfers at gigabit per second rates for backing up databases or migrating virtual machines, and the network must be equally adept at supporting fast queries among many sites where the data is stored.

Figure 2 shows a typical architecture for web-based applications. Typically application-layer SSL/TLS encryption and authentication protects front-end traffic between the browser and the server. Remember that sensitive traffic exchanged among back-end servers and databases is also vulnerable and must be protected. While VLAN segmentation is often used to separate traffic in the back-end network, it does not provide sufficient protection for sensitive information in an environment that is shared among many customers. Encryption and authentication on a frame-by-frame basis provides privacy and data integrity that meet regulatory compliance mandates and follow data security best practices. Network traffic among back-end components must be protected without slowing down applications and while allowing servers to migrate and scale up and down with demand.



**Figure 2: Front-end vs. Back-end processing in typical web applications**

Depending on the architecture, the back-end network that supports the web servers may consist of three separate networks (see Figure 3): the cloud network, the wide area network (WAN), and the data center local area network (LAN). By establishing secure connectivity among servers in the cloud and the data center using these networks, we can build a flexible and scalable cloud architecture that minimizes cost while supporting business-critical applications and services. When we want to add new applications and services, this architecture allows us to add new applications where they fit best: in the data center or in the cloud environment.



**Figure 3: Components of Back-End Networks**

It is important that the network is flexible and it should provide cost-effective connectivity. For the WAN, private networks provide a degree of consistency and

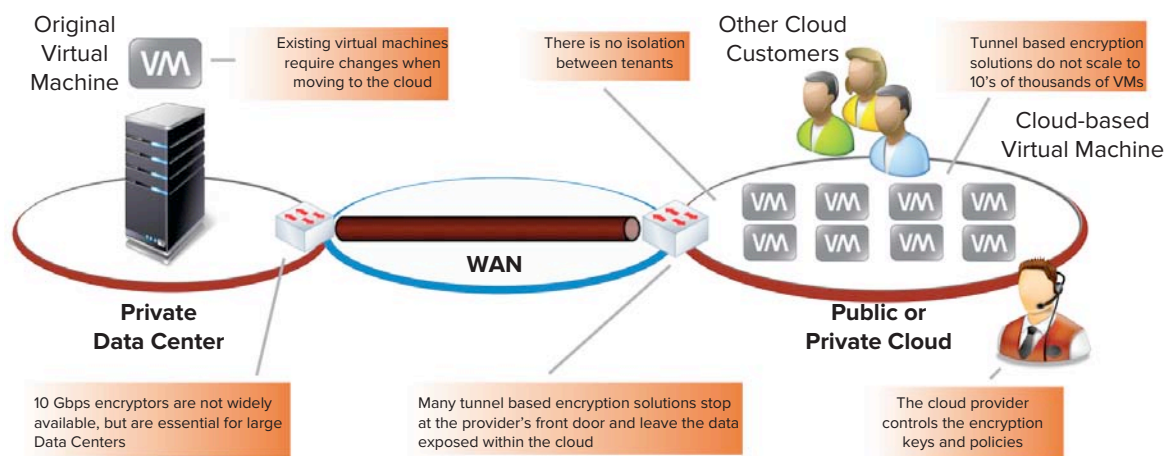
reliability that the public Internet cannot match. Private Layer-3 networks and MPLS networks are good options, but Layer-2 Ethernet is a compelling solution for connecting data centers and cloud environments because the cost of bandwidth is very low compared to other options and because the network is flat and simple to manage. Many organizations already have Layer-2 Ethernet connectivity among several locations, and it's usually easy to add a connection at the cloud provider's site. It's also important to consider scalability in designing the back-end interconnectivity solution.

### Summary of architectural principles

- Run applications in the environment in which they are best suited to run
- Use dedicated servers where appropriate and virtualize everything else
- Take advantage of cloud efficiencies, while leveraging existing data centers and in-house virtualization
- Design the back-end network with as much care as the front-end
- Use private networks for interconnecting data centers and cloud environments, and consider that Layer-2 Ethernet may be the best option
- Encryption is essential for data privacy in the back-end network
- Consider flexibility, network usage costs, latency and throughput performance, scalability and security when planning connectivity among back-end resources

### Common problems

Now let's discuss some of the issues that come up when trying to secure data in motion among servers in a cloud environment or mixed cloud and data center environment. Figure 4 illustrates some of the issues.



**Figure 4: Common problems encountered when securing network traffic**

### Data privacy and integrity

While VLAN tagging or MPLS labels can provide traffic separation on a shared network, neither offers the privacy afforded by encryption. Regulatory requirements often require encryption in a shared network. While legacy point-to-point VPN tunnels are sometimes used to provide confidentiality and integrity, these solutions usually fall short of expectations in terms of performance, scalability and management costs. Existing solutions also may leave gaps where the traffic is not protected. Usually traffic is decrypted at the cloud provider's front door (where network traffic enters the cloud provider network). This leaves your traffic exposed to network sniffing and traffic injection throughout the cloud provider's network.

In *Security Guidance for Critical Areas of Focus in Cloud Computing V2.1*, the Cloud Security Alliance provides the following guidance:

*Although cloud provider networks may be more secure than the open Internet, they are by their very architecture made up of many disparate components, and disparate organizations share the cloud. **Therefore it is important to protect this sensitive and regulated information in transit even within the cloud provider's network.***

In the 3.0 version of the document, the Cloud Security Alliance provides the following guidance:

*Encrypt using sufficiently durable encryption strengths (such as AES-256) that comply with the same corporate and regulatory mandates used for encrypting internally maintained files.*

A recent report from the U.S. Office of the National Counterintelligence Executive points out the risks of cloud computing:

*Although cloud computing offers some security advantages, such as robust backup in the event of a systems disruption, the movement of data among multiple locations will increase the opportunities for theft or manipulation by malicious actors.*

"Foreign Spies Stealing U.S. Economic Secrets in Cyberspace", October, 2011

### Isolation from other cloud customers

Without isolation, other customers can attack your assets in the shared cloud environment. This happens when malicious cloud customers attack other cloud customers. This can occur when an attacker compromises the assets of a legitimate cloud customer and uses these compromised assets as a jumping off point to attack other cloud customers. Cryptographic isolation is best because it

uses strong cryptographic authentication techniques to block unauthorized traffic.

### **Performance**

The bandwidth of the data center connection must be sufficient to support high bandwidth applications such as VM migration and database backups that can total hundreds of gigabytes each. Data centers often require encryption speeds in the range of 1 – 10 Gbps in order to connect to other data centers and resources in the cloud. A minimal delay through the network is also essential to support latency-sensitive applications and to maximize the performance of back-end processing systems.

Encryption should be used to protect the traffic entering and leaving the data center where network traffic is often concentrated. It is important that the encryption solution can keep up with this potentially heavy traffic load, and typically dedicated hardware is required. Many existing solutions cannot encrypt traffic at 10 Gbps with low latency, so they drop or delay traffic and application performance suffers.

While encrypting traffic at a few points works well for the private data center, it is not ideal for the cloud environment. In the cloud environment, we'd like to be able to terminate the encryption on each server that uses the data. Tunnel-based solutions typically have a single tunnel that terminates at one point in the cloud (because of inherent scalability and management issues with tunnels). Unfortunately one tunnel server may not have enough CPU capacity to encrypt traffic from many other servers, and it may become a performance bottleneck. A better approach is to distribute the encryption workload across as many physical servers as possible.

### **Scalability**

As the number of servers that are to be interconnected increases, many solutions cannot scale sufficiently to provide full mesh connectivity. Solutions that are tunnel-based often have limited scalability or require multiple network hops that force the traffic through specific sub-optimal paths in the network (hairpinning) and require the traffic to be decrypted and re-encrypted at each hop, which further degrades performance and introduces additional latency.

### **Redundancy**

Tunnel-based point-to-point VPN solutions often fail to provide hitless redundancy. With tunnel-based VPNs, only one point on the other side of the network has the keys to decrypt the traffic. This means that when there is a failure of a tunnel endpoint, the remaining tunnel endpoint must detect that the tunnel is down and then establish a new tunnel or switch over to a backup tunnel. Traffic is lost during this time.

### **Controlling the encryption keys**

Few network encryption solutions allow you to control your own encryption keys. In order to protect your sensitive information, you must control the keys. Otherwise it's like locking your front door, and leaving the key in the lock. Your information security auditor may require you to demonstrate that you control the encryption keys. If the cloud provider possesses your encryption keys, he may be forced to decrypt your data and provide it to law enforcement authorities. This also makes it possible for malicious or disgruntled employees of the cloud provider to steal your data. *Security Guidance for Critical Areas of Focus in Cloud Computing V3.0* is a document authored by the Cloud Security Alliance (CSA). It states the following requirement:

*In order to maintain best practices and pass audits the organization should **manage their keys in the custody of their own enterprise** or that of a credible service from a cryptographic service provider.*

### **Limited flexibility**

Many solutions force you to encrypt network traffic at Layer-3. This may work for some applications, but it can also cause problems. Encrypted tunnels can be used to bridge between Layer-2 subnets, but tunnels introduce scalability issues and management complexity. On the other hand, encrypting with per-packet authentication at Layer-2 without tunnels allows full mesh connectivity among all of the devices and virtual machines on the subnet, and it allows virtual machines to move around the environment without changing IP addresses.

Existing solutions often limit flexibility because they do not provide fine-grained control of the encryption policies that specify which traffic is encrypted, dropped, or passed in the clear. These solutions may be adequate for remote access VPNs or site-to-site VPNs, but they don't have the powerful policy controls that allow applications and services to be segmented or isolated according to policy controls.

### **Changes to existing applications**

If the cloud network is not bridged at Layer-2 to a subnet of the data center, the IP addresses of virtual machines must be changed when the virtual machines run in the cloud environment. This can be disruptive and difficult to manage and may require changes to applications and firewall rules. Frequent firewall changes increase complexity and introduce the potential for errors and connectivity issues.

### **Summary**

The need to securely connect resources in data centers and cloud environments is clear, but existing solutions have significant drawbacks. Furthermore, most existing solutions are not flexible enough to support the changes that will

inevitably occur in the next few years as cloud and data center environments evolve. In the following sections we will examine some of the characteristics of next-generation solutions that will allow your solution to adapt, scale and take advantage of new connectivity options.

### **Group Keying: The “Key” to Cloud Security**

The characteristics of cloud environments typically include the terms “scalable”, “dynamic” and “elastic” – terms that are nearly opposite of the user experience with IPsec tunnels established with IKE (the predominant method of securing connections to virtualized private data centers and public cloud infrastructures). The IKE/IPsec method also leaves control of the policies and keys with the provider, which violates data privacy regulations. Furthermore, tunnels force point-to-point connectivity that is limited to two endpoints per tunnel. This point-to-point mapping of large-scale distributed networks is exponentially complex to the point of being prohibitive in terms of both the cost to manage and the performance limitations it imposes.

Group keying, however, is ideally suited for the cloud environment due to its elegant scalability, easy management and its ability to allow policies and keys to be controlled centrally, from the client side of the secure connection. Group keying eliminates point-to-point negotiation of keys, which becomes intractable as the number of endpoints grows.

With group keying, keys are generated centrally and then securely distributed to authenticated group members. Once the keys are distributed to the group, any group member can communicate securely with any other member. The group key used to encrypt data destined for one group member is the same key used to encrypt data destined for all of the group members. This is very different from tunnel-based solutions. With tunnels, the system maintains a unique key for each endpoint. Maintaining this information and looking up which key to use for every incoming packet degrades the performance and limits the scalability of tunnel-based solutions. Furthermore, managing these point-to-point tunnel-based solutions and adding additional tunnels becomes difficult and time consuming after a handful of tunnels are provisioned.

### **Essential solution components**

It’s clear that virtual machines (VMs) executing in a cloud environment need to be able to communicate securely with VMs and physical servers in a private data center; however, secure connectivity is not enough. The following are essential solution components for securing back-end network traffic and solving the issues described in the previous section.

### ***Encryption and authentication***

The only way to provide data privacy, data integrity and cryptographic isolation is to encrypt and authenticate each network packet or frame. Security best practices require the use of NIST-approved encryption and hashing algorithms that are validated and tested according to the U.S. government FIPS 140-2 specifications under the Cryptographic Module Validation Program (CMVP).

### ***Cryptographic Isolation***

Traffic should be encrypted and authenticated (per frame) such that it is unreadable and unusable to all other cloud customers who may receive the traffic accidentally or through illicit means. Cryptographic isolation also protects the data center from threats originating in the cloud by blocking all traffic that is not authenticated.

### ***Performance: high bandwidth and low latency***

Having the option to choose between hardware or software encryption on a per-site basis allows you to determine the appropriate balance between performance, cost and user self-provisioning. You may choose to deploy hardware-based encryption only where is required to meet critical performance and latency objectives and only if it does not interfere with user self-provisioning. Software-based encryption that is distributed among servers can be used everywhere else.

Solutions that have dedicated hardware that is optimized to perform cryptographic functions can keep up with line rate traffic at 10 Gbps while providing low latency. Because data centers often have a small number of large network pipes connecting them to a wide area network, hardware-based encryption is often essential for encrypting traffic to and from data centers. Hardware encryption appliances need to be deployed in the network in advance, so they are often best suited to data centers that have predictable traffic patterns.

Software-based encryption appliances can be allocated on-demand and distributed throughout the virtualized environment. Tunnel-based solutions often concentrate traffic to a single point and try to encrypt all of the traffic on one server. The problem with this approach is that the server can become a bottleneck to performance. Group keying solutions distribute the encryption workload throughout the environment and encrypt and decrypt the traffic on the servers that use the data. This protects the data more effectively and avoids bottlenecks by distributing the encryption workload across the servers. With group keying, encryption capabilities expand as the workload expands, so capacity scales to match the workload. This is exactly how cloud services should work.

### **Scalability**

Gartner analysts Thomas J. Bittman and Lydia Leong estimate that the number of

virtual machines deployed in IaaS (Infrastructure as a Service) service provider offerings will nearly double each year through 2014. As the number of virtual servers in the environment grows, scalability of the security solution will become increasingly important. The group keying technology pioneered by Certes Networks provides a solid foundation for encrypted full-mesh connectivity among tens of thousands of endpoints. This proven technology has been deployed and working for years in hundreds of networks throughout the world. While group keying is important for networks containing hundreds of network nodes, it is absolutely essential to allow cloud environments to scale to tens of thousands of servers.

### ***Redundancy***

Group keying can provide hitless redundancy of encrypted traffic. Both redundant decryptors have the same decryption key, so there is no need to load a new key or establish a new session when a failure affects one of two redundant decryptors – the other decryptor is always ready to receive traffic from any other group member. Group keying also allows load balancing to work seamlessly. Simply put, group keying allows the network to work the way you designed it in the first place.

### ***Customer control of the encryption keys***

Cloud customers are often required to maintain control of their own policies and keys for regulatory reasons. Likewise, cloud providers may not wish to bear the financial and legal burdens associated with being in possession of the keys. You should not settle for a solution that does not allow you to control your own encryption keys.

### ***Flexibility***

Solutions that are policy-based and allow you to simply and easily select which traffic to encrypt, drop or pass in the clear have a definite advantage. These capabilities give you the control you need to protect the network and limit connectivity among applications and services. Many tunnel-based solutions do not offer full featured policy controls that allow you to specify the connectivity and encryption policies that you need.

Solutions that let you choose what part of the packet to encrypt (this is known as multi-layer encryption) give you the ability to protect many different types of networks with the same solution. This allows you to choose the network connectivity that best fits your application, budget and availability from your network and cloud provider. With multi-layer encryption, you can encrypt nearly any type of network, including:

- Private Layer-2 Ethernet

- Public or private Layer-3 IP
- MPLS

As high bandwidth and low latency Ethernet connectivity becomes more widely available, it will likely become the most cost-effective option for many organizations to connect existing data centers to VMs in the cloud. Layer-2 network connectivity coupled with Layer-2 encryption using group keying provides any-to-any connectivity without tunnels. While Layer-3 (IP) networks can tunnel Layer-2 traffic, only Layer-2 (Ethernet) networks can support Layer-2 connectivity without tunnels.

As your needs change or as new options or lower cost alternatives are offered through your carrier, a solution that supports multi-layer encryption with powerful policy controls allows you to migrate to and protect the new network without changing your encryption solution.

### **Changes to existing applications**

If cloud networks are bridged to data center networks, virtual machines do not need to change their IP addresses when moving from the data center to the cloud. Bridging allows the same subnet to be used in the data center and in the cloud, so hosts or servers cannot distinguish between the two environments – hosts or servers work the same way in both environments without changes. While this type of bridging is possible today, it is not practical without authentication.

Authentication allows you to bridge subnets in the data center with subnets in the cloud network with confidence. Authentication protects machines on a trusted network from unauthorized machines that would try to attack them. It goes beyond the protection offered by firewalls by using cryptographic signing techniques. Group authentication is based on a secret key that is securely distributed to all of the group members. The key is distributed only to the group members that you designate in your security policy, so you control group membership. Group members check the authentication code of each incoming packet to confirm the sender signed the packet using the same key, and packets that are unsigned or signed with the wrong key are dropped.

Based on the security foundation provided by authentication, you can use bridging to connect subnets among the data center and the cloud. The result is that virtual machines can migrate among data centers and cloud environments without changing IP addresses and without changing firewall rules.

## Conclusion

The technology that enables these essential solution components has been proven over nearly ten years of deployments to government agencies, defense contractors, financial organizations, as well as enterprises looking to secure PCI, HIPAA and other compliance-driven sensitive information. Certes Networks is building on these proven technologies to provide a solution that makes the cloud safe for sensitive workloads. This will lead to secure and optimized solution architectures for data centers and cloud environments that dramatically reduce IT costs. You can read more about Certes Networks technology, products and solutions at [www.CertesNetworks.com](http://www.CertesNetworks.com).

## About Certes Networks

Certes Networks is the leader in multi-layer encryption solutions for high performance networks. The Company provides advanced IPSec VPN and encryption solutions for wide area networks, and enables secure connectivity to private and public clouds. Certes Networks helps organizations improve security, decrease risk, and reduce the cost of compliance while enabling high performance and secure connectivity to critical infrastructures.