White Paper

Integrating the Cloud: Bridges, Brokers, and Gateways

Organizations are increasingly adopting a hybrid strategy for cloud computing to realize the benefits without compromising on control. The integration at the network, infrastructure, and process layers that is required by hybrid models can be addressed by three new cloud-focused solution types: bridges, brokers, and gateways.

by Lori MacVittie
Senior Technical Marketing Manager

Contents

Introduction

For many organizations, a hybrid cloud architecture is the best way to take advantage of the benefits of public cloud computing resources without compromising on security or operational control. Hybrid cloud architectures can be enabled in several ways, each offering various degrees of operational consistency and control over the integrated environment. As adoption rates of hybrid architectures increase, solutions designed specifically to address challenges with implementation will continue to emerge.

These challenges have a common theme: integration. Whether the concern is reining in control of applications and processes for Software as a Service (SaaS), extending the network beyond the data center into Infrastructure as a Service (IaaS), or managing cloud storage services, the key to success is integration.

> *At present, two-thirds of mid-sized firms indicate that their cloud solutions are only minimally integrated into their IT environments and function independently of other IT operations. In the future, these users expect this situation to change; over one quarter expect full integration into their IT environments, and an additional 20 percent expect a major shift to cloud-based solutions (suggesting a higher level of integration) as the use of on-premises solutions declines.*
>
> *"Cloud Computing in the Midmarket: Assessing the Options," IDC, September 2011*

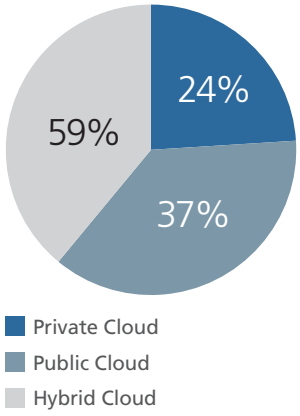

Source: 2011 Future of Cloud Computing Survey – June 22, 2011

Figure 1: Hybrid cloud implementations are the largest, and still growing, segment.

Integration has long been the bane of those within IT who are focused on applications, but that long struggle has produced best practices and guiding principles that apply to more network-oriented integration efforts as well. Applying these principles leads to three categories of solutions focused on integrating data centers with cloud computing resources: bridges, brokers, and gateways.

Each of these solution categories solves specific issues with integration across the full spectrum of network, compute, and storage stacks. Which solution is best suited for a given organization depends upon several factors, including:

- The level of operational consistency required.
- The level of automation and orchestration desired.
- The importance of control over all components, both on and off the premises.

Regardless, careful consideration and understanding of these solutions is required to ensure the choice of an architectural approach that's appropriate to meeting requirements and expectations.

Bridges

Wherever two networks exist, someone will try to connect them. Cloud bridges address network connectivity between cloud and data center networks, enabling a secure connection over which data can be transferred and resources managed.

Bridges enable security not only by virtue of using IPsec but by controlling access to resources through a strategic point of control in the network. By tunneling management functions through a bridge, access to and governance of remote resources can be more effectively managed with existing enterprise systems.

There are two types of cloud bridges: asymmetric and symmetric. Both take advantage of IPsec to secure a tunnel between the data center and the cloud; the difference lies in whether end-points are similar or dissimilar.

Asymmetric Bridges

An asymmetric bridge is one in which the IPsec endpoints are dissimilar. For example, an asymmetric bridge can be deployed using one IPsec provider in the data center connecting to Amazon's Virtual Private Cloud (VPC) IPsec endpoint.

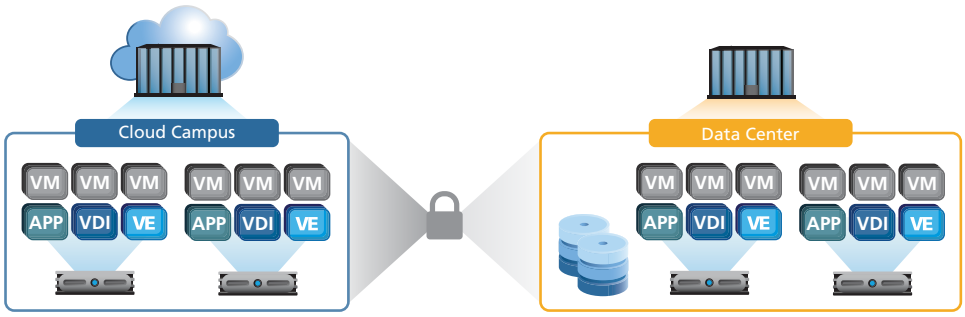

Figure 2: In an asymmetric bridge, the cloud IPsec endpoint is not the same IPsec endpoint as in the data center.

An asymmetric bridging solution reduces operational costs by eliminating the need to manage one-half of the solution remotely, but this also brings a disadvantage: a lack of control over the remote endpoint, resulting in a loss of visibility that can complicate troubleshooting.

Symmetric Bridges

A symmetric bridge is one in which the IPsec endpoints are deployed using the same provider.

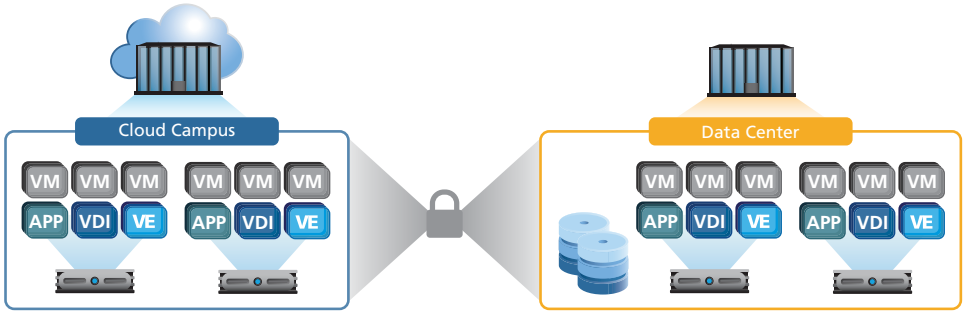

Figure 3: A symmetric bridge establishes an IPsec tunnel between the data center and the cloud, using the same technology.

Using the same provider at both ends of the connection has the advantage of operational consistency and offers more control over access as well as traffic flows. Depending on the IPsec solution, this model may also enable additional value by deploying complementary delivery services such as WAN or web optimization services.

The disadvantage is that it necessarily requires the bridging solution to be available for deployment in cloud environments, generally in a virtual or software form-factor. While availability of such solutions is growing, limitations still exist that may constrain the organization's ability to choose a cloud provider.

Network Normalization

In addition to bridging two networks securely, these connections can also serve to bridge two environments by normalizing network communications, much as a bridging switch does within the data center. In general, bridges accomplish this by creating network overlays.

A network overlay is also referred to as network virtualization, a technique in which a logical network topology is layered over the existing physical and logical layer 2 and layer 3 network. These network overlays provide the means to route traffic between the two domains primarily by identifying traffic as part of a "virtual network" using customer-specific identifiers embedded in extended Ethernet frames. Protocols

enabling this normalization include emerging specifications such as VXLAN and NVGRE as well as established but lesser-known standards like IP in IP (RFC 2003).

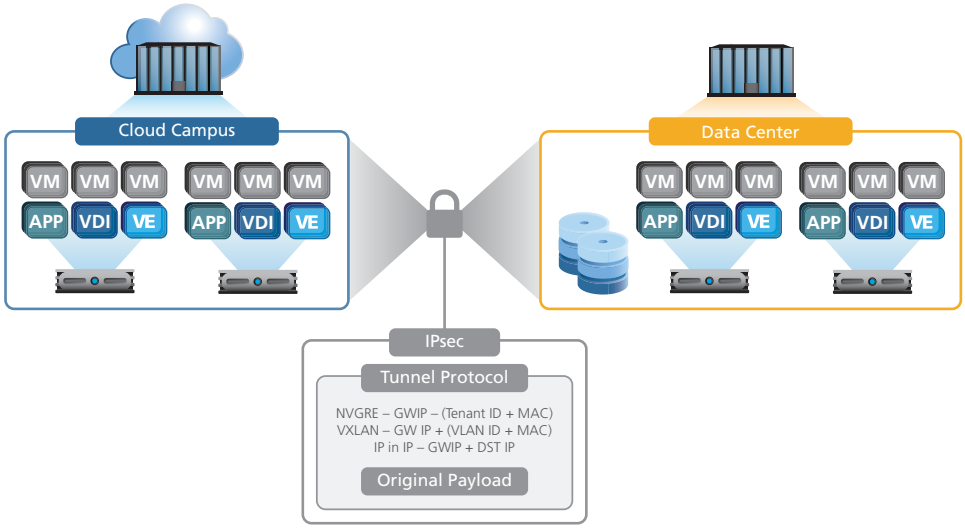

Figure 4: Network connectivity is normalized and secured between environments using a variety of protocols.

Brokers

Cloud brokers integrate cloud and data center environments at the process layer. For example, an identity broker mediates authentication and authorization processes for cloud-deployed applications and services as a means to centralize and maintain control over access and accounts. This mediation entails several steps and requires integration between systems at relevant layers of the stack.

Brokers, which can enforce a variety of application delivery-related policies, are an architectural remedy to the challenge of managing distributed applications, particularly when those applications are deployed in environments over which IT has virtually no visibility or control. SaaS is a prime example. By mediating access through the IT-controlled infrastructure, corporate policies can be enforced without sacrificing the benefits associated with a public cloud deployment.

The most common broker patterns are those governing identity, access, and delivery.

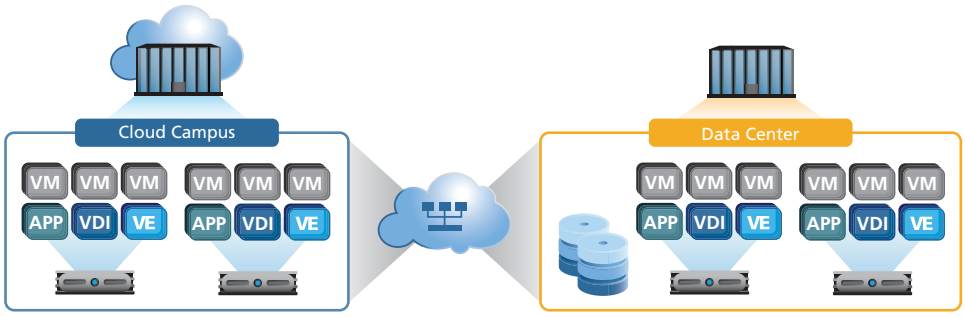

Figure 5: Cloud brokers integrate at the process layer to mediate and enforce application delivery policies.

Cloud Identity Brokers

A cloud identity broker mediates authentication and authorization between the data center's identity management systems and cloud-deployed applications, which generally maintain their own identity stores. The purpose of a cloud identity broker is to ensure access only by authorized users in real time, a role made necessary by incomplete integration between remote applications and local identity stores.

Depending on the cloud application for which credential management is being brokered, the application may or may not maintain its own identity store. SaaS applications undoubtedly do so, while custom applications may be written specifically to use assertion tokens instead. Such tokens may take multiple forms, and applications may be further secured by only accepting brokered connections through authorized data center delivery systems.

Where remote identity stores are in use, the synchronization that is often performed manually can lag behind changes in authorization, leading to orphaned accounts that grant access to users or employees after they have separated from the organization. A cloud identity broker forces authentication and authorization through the authoritative systems within the data center, ensuring credentials are up to date. The cloud-hosted application, configured to trust the data center's authorization, honors its assertions.

While a cloud identity broker does not eliminate the need for manual synchronization, it does eliminate many of the associated risks.

Cloud Access Brokers

Despite the marketing hype touting unfettered access to cloud applications as a benefit, access control is required for a variety of reasons—many of them related to industry and government regulations. Access management services, however, are not part and parcel of most cloud offerings and thus must be "bolted on" in some way to ensure access is granted only to those authorized.

Much like cloud identity brokers, cloud access brokers mediate on behalf of users with cloud applications and allow or deny access to specific resources based on corporate policies. Depending on the implementation, such policies may be loosely defined (at the application level) or granular in nature (specific URIs, for example).

The technical implementation of cloud access brokers can be achieved by two primary architectural options:

- **Full-proxy broker:** All requests are tunneled through and managed by the broker. This model can be implemented a number of ways, including through the use of a tunnel (i.e., a VPN model) or by virtualizing the service.
- **Half-proxy broker:** Requests are validated by the broker, but subsequent communications occur directly between the client and the service. This model is based on a delayed or late binding model in which an intermediary determines how requests should be handled and then removes itself from the exchange.

Cloud Delivery Brokers

Cloud delivery brokers mediate on behalf of the user to ensure that availability and performance requirements are met. A cloud delivery broker implementation relies heavily on context, taking into consideration variables such as location, business requirements, cost, and security. The cloud delivery broker uses this information to make intelligent decisions regarding the delivery of applications to users as a means to meet service-level agreements (SLAs) as well as to avoid service disruptions.

A cloud delivery broker can assist in making decisions both as to where users are directed upon requesting resources and where resources are best deployed to meet business and operational SLAs. To achieve SLA goals, applications or their owners must be able to specify constraints such as:

- Location restrictions.
- Accessibility limitations (e.g., cannot run in public clouds, must run in public clouds, etc.).

- Latency and bandwidth requirements.
- Performance expectations.
- Cost limitations.

Conversely, both public and private cloud environments must be able to describe capabilities, such as:

- Bandwidth.
- Location.
- Cost.
- Type of environment (public, private, community).

Using these two sets of characteristics, a cloud delivery broker can determine the environments available to it that are best able to meet the requirements both during deployment of the application instance and when access to the application is requested.

Cloud balancing and business continuity initiatives require cloud delivery brokers, as they are the components that ultimately make cloud-level routing decisions for users requesting access.

Gateways

Traditionally, gateways in data centers are a transition point from one protocol to another, such as from IPv6 to IPv4, or from one network to another, such as forward proxies that transition traffic from internal networks to the Internet. Cloud gateways are similar in that they create transitions between environments by integrating cloud and data center environments at the resource management layer via APIs.

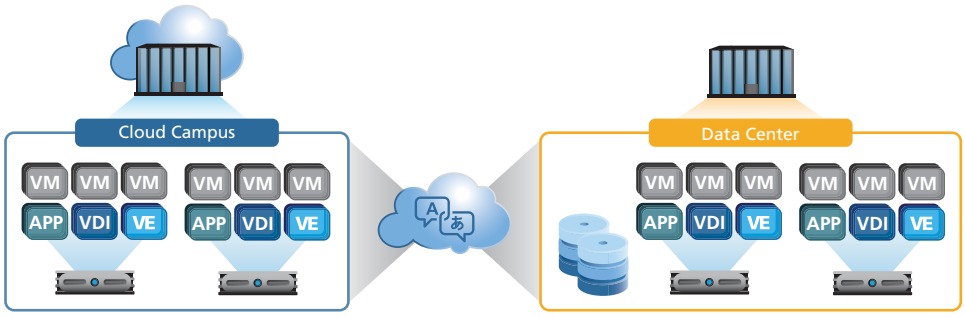

Figure 6: A cloud gateway creates transitions between environments by integrating resources through APIs.

A cloud gateway may be specific to a resource type, such as cloud storage gateways, or it may be more generic, integrating multiple resource types by transitioning through cloud provider API frameworks.

The key to a cloud gateway is the ability to integrate cloud-deployed resources into the data center's architecture so as to consistently apply security, performance, and availability services to meet business and operational requirements.

Storage Gateways

Cloud storage gateways were the first cloud gateways to emerge. These solutions integrated cloud-deployed storage services with data center storage networks to provide automated tiering and support backup strategies. The exponential growth of storage in the data center necessitates that costs be constrained without compromising the organization's ability to comply with regulations regarding retention and security of data. Cloud storage services enable IT to fulfill these requirements.

Cloud storage gateways not only enable extension of corporate storage networks into the cloud, but often also provide additional functionality around backup and tiering procedures that can make more efficient use of cloud and corporate storage.

Cloud Gateways

Cloud gateways are elusive creatures, with very few true examples in existence today. Due to historical definitions, the term is often used to describe everything from simple bridging solutions to more complex intercloud management systems. While cloud gateways may comprise cloud bridging functionality, a gateway must provide more than simple network connectivity to enable the extension of the data center and the integration of cloud-deployed resources into the data center.

A cloud gateway should necessarily interact with cloud framework APIs as a means to integrate remote resources into the management and operational architectures of the data center. This model realizes operational consistency and compliance with business and operational requirements—including those related to security, performance, and availability—that cannot necessarily be achieved when two separate sets of processes and resources must be coordinated manually.

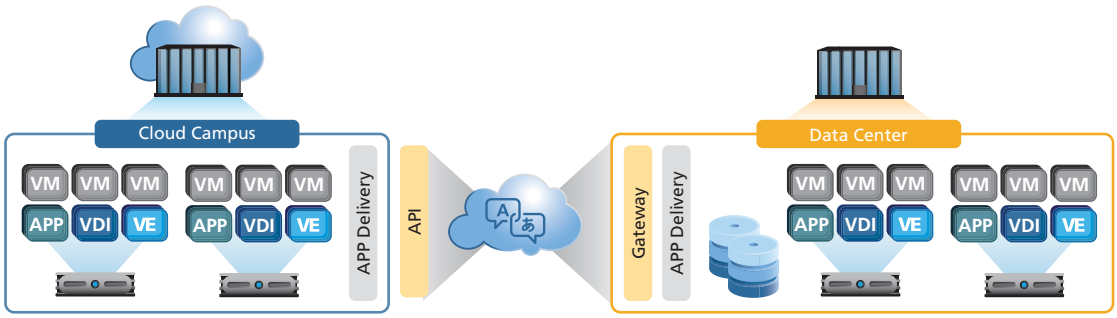

Figure 7: A cloud gateway communicates through APIs to enable integration of cloud-deployed resources into existing data center architectures.

A cloud gateway can further insulate organizations from potential lock-in by ensuring the ability to move from one provider to another simply by directing the gateway at a different cloud.

Conclusion

The benefits of public and hybrid cloud computing are compelling enough to engender interest from and adoption by organizations both large and small. And yet the challenges and concerns over how to integrate with cloud environments and maintain control over security and performance continue to inhibit broader adoption.

Bridges, brokers, and gateways are both architectural and product-based solutions. As demand for integration continues to grow, so will the range of solutions available to organizations desiring to leverage them to extend their data centers into the cloud. Recognizing these new breeds of solutions will enable organizations to choose those that best suit their unique integration needs and plans for cloud computing in the future.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com