embrane
Powering the Agile Network

Powering Virtual Network Services

Executive Summary

Cloud computing has the potential to dramatically alter the ways in which IT resources are deployed and consumed. The advent of virtualization technology has revolutionized the server and storage industries and laid the foundation for today's cloud computing initiatives. Unfortunately, the network has now become a major impediment to cloud adoption. Conventional networking solutions weren't designed for the cloud and aren't well suited for addressing the new design requirements. When it comes to layer 4-7 network services, neither hardware-based appliances nor virtual appliances can provide the required agility to support the new infrastructure paradigm.

Embrane offers the industry's first network services platform specifically designed for the cloud. Embrane's distributed architecture enables the rapid and efficient delivery of highly extensible and scalable layer 4-7 on-demand network services. This white paper provides an overview of the Embrane Virtual Network Services Platform and explains how Embrane helps cloud service providers and enterprises deploying on-premise private clouds deliver highly scalable and elastic network services on-demand.

Cloud Networking Challenges

Enterprises are turning to the cloud – either public, private or a hybrid model – to improve agility, reduce expenses and risks, and accelerate innovation. Cloud computing redefines the way IT assets are deployed and utilized, and impacts the way all layers of the network are architected and managed. When it comes to layers 4-7 in particular, conventional enterprise network services appliances – firewalls, load balancers, VPN devices, etc. – aren't well-suited for cloud deployment models. Service providers implementing public clouds and enterprises deploying on-premise private clouds must identify more scalable, agile and cost-effective ways to enable virtualized, on-demand network services.

IT organizations rely on a variety of networking solutions to safeguard IT assets, ensure efficient use of resources, and optimize service quality. Firewalls, VPNs and intrusion detection and prevention solutions help mitigate security risks and improve network availability and integrity, and devices like load balancers and WAN optimization controllers help boost application performance and improve end-user experiences. Many innovative cloud providers are looking to deliver such value-added layer 4-7 networking functions in the cloud in the form of on-demand services.

Until now instantiating network services in the cloud has been a challenge. Most layer 4-7 network appliances were intended for traditional enterprise use. They weren't designed to support multiple tenants or on-demand or pay-as-you-use services and aren't easily adapted for cloud deployment models. As a result, many cloud providers simply don't offer value-added network services. Instead they let end-customers deploy and administer networking appliances on their own, very often by simply allowing them to run a virtual appliance inside one of the virtual machines already provided as part of the generic cloud service[1]. They miss out on incremental service revenue and competitive differentiation, and they risk lost sales, prolonged implementation cycles, and customer dissatisfaction.

Figure 1 – Over-the-Cloud vs. Under-the-Cloud Deployment Options

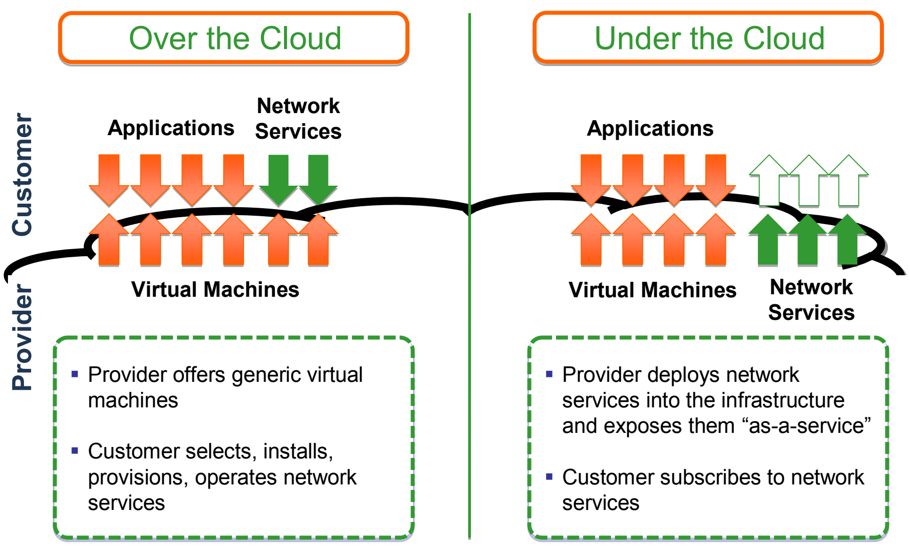

[1]This approach is sometimes referred to as an "over-the-cloud" deployment model.

Forward-looking providers are looking to offer layer 4-7 network services in the cloud to accelerate customer deployments, move up the value chain, and enable new revenue streams. But delivering on-demand network services in an efficient and profitable manner can be a challenge. Providers have two options for deploying "under-the-cloud" network services today: dedicated hardware appliances or software-based, virtual appliances. Each of these approaches offers distinct advantages and disadvantages. Neither is particularly well suited for the cloud.

Hardware-based appliances offer high performance and reliability, but are expensive, inefficient and inflexible. Originally designed for the enterprise, most hardware appliances were not intended to support multi-tenancy. They can't be logically partitioned to support multiple customers without sacrificing performance-guarantees or restricting features. To even resemble a multi-tenant environment, service providers often have to over-provision infrastructure, deploying dedicated devices for each customer and each network service. Implementing and provisioning services is usually a non-automated, time-consuming task, which often takes days to complete. The drawn out process drastically impacts the service provider's ability to provide a timely response to the customer, and renders the service provider unable to offer dynamic, pay-as-you-use services.

Software-based, virtual appliances are easier to provision and less CAPEX-intensive than dedicated hardware solutions, but they pose capacity planning and network design challenges. Virtualized appliances run on industry-standard x86 servers, but in order to ensure that the appliance delivers the desired level of performance, these servers have to be properly sized to meet the anticipated capacity and performance requirements of a particular application or customer; similarly, the virtual compute resources (vCPU, vRAM, etc) allocated to the appliance need to be properly sized. In many cases, virtual appliances have to be re-hosted over time to address increasing or decreasing capacity demands. Moving a virtual appliance from one machine to another in a seamless fashion can be a challenge.

Whether hardware-based or software-based, conventional layer 4-7 network appliances simply aren't well-suited for delivering cloud-based services. Each appliance – whether virtual or physical - represents a discrete management point. As service providers add capacity or customers, management complexity grows and operations expenses increase. Worse still, most conventional networking appliances were originally conceived for the enterprise and aren't inherently architected to support multi-tenancy, service elasticity, or pay-as-you-use service models. In addition most networking appliances provide very limited programming interfaces that weren't designed to support sophisticated service orchestration tools or other Operations Support Systems (OSSs).

Figure 2 – Conventional Networking Appliances aren't suited for the Cloud

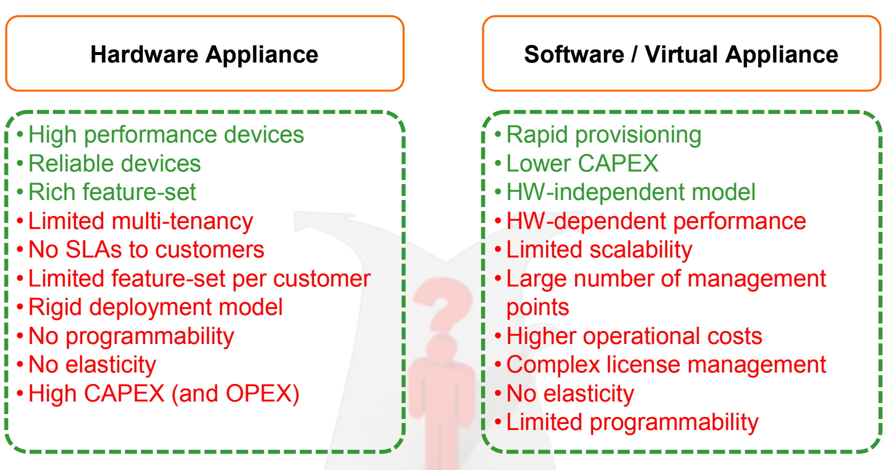

Enabling Network Services Under the Cloud

Embrane offers the industry's first layer 4-7 network services platform specifically designed for the cloud. Embrane's Virtual Network Services Platform combines the performance, scalability and reliability advantages of a hardware-based appliance with the flexibility and simplicity of a virtualized, software-based solution. Better still, it is the first platform deliberately designed to enable "under-the-cloud" network services by leveraging a fully distributed architecture to support fundamental service provider requirements like multi-tenancy, elasticity, and a true usage-based pricing model.

Virtualization technology has revolutionized the server and storage industries and laid the foundation for today's cloud computing initiatives. Using the latest server and storage virtualization solutions IT administrators can efficiently provision compute and storage resources in real-time to enable on-demand IT services. But provisioning the accompanying network services can be a resource-intensive, time-consuming proposition. Embrane's Virtual Network Services Platform is the first solution to bring the full benefits of virtualization to networking. The platform eliminates implementation obstacles and accelerates service deployment, allowing cloud administrators to efficiently enable layer 4-7 network services on-demand.

A software-based solution, the Embrane Virtual Network Services Platform is implemented using industry-standard x86 servers and exploits hypervisor technology to eliminate vendor lock-in and contain costs. The platform leverages a distributed virtual appliance architecture that decouples network services functionality from the underlying physical infrastructure and hypervisor technology, providing high scalability, flexibility and performance.

Embrane's Virtual Network Services Platform Raises the Bar on Agility and Scalability

The Embrane Virtual Network Services Platform was designed from the ground up to enable cloud-based network services. The platform includes two distinct components for provisioning and enabling agile layer 4-7 network services under the cloud: Distributed Virtual Appliances (DVAs) and the Elastic Services Manager (ESM).

Figure 3 – Embrane Virtual Network Services Architecture Overview

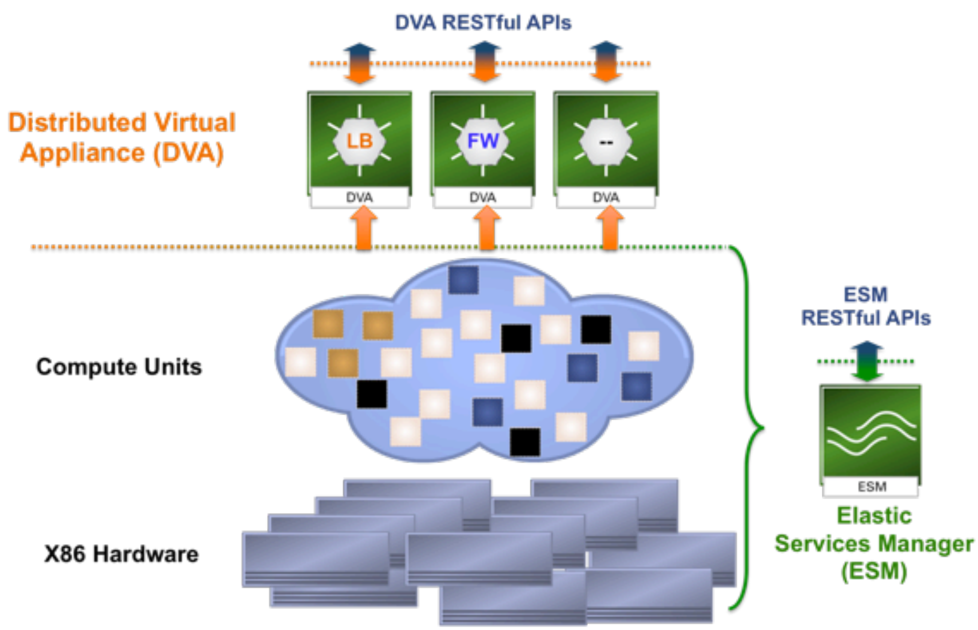

Distributed Virtual Appliances

DVAs are logical network service containers that are instantiated across a pool of general-purpose x86 servers. Current examples of DVAs include firewalls, load balancers and IPSec VPN appliances. DVAs are essential for enabling multi-tenancy, elasticity and on-demand service delivery.

A DVA is implemented by aggregating the compute power of several virtual compute resources, called Compute Units. A Compute Unit (CU) is a VM supported by any type of hypervisor technology. The CU provides granular access to x86 compute resources (a core or fraction of a core) and gives cloud providers the freedom to decouple their choice of x86 server hardware and virtualization technology from the implementation of network services. CUs provide an abstraction layer above the physical hardware that shields network and IT administrators from underlying environmental complexities.

A DVA is an abstract network resource that is realized across a cluster of Compute Units. A logical network element with a single point of management, a DVA allows a cluster of CUs to be administered as a unified entity. DVAs can span layer 2 and layer 3 domains, leveraging Compute Units located anywhere in the data center, independent of network topology.

DVAs are the key to delivering truly elastic services and eliminating capacity planning complexities. IT teams can dynamically reconfigure DVAs or throttle DVA capacity in real-time by simply adding or removing Compute Units in a non-disruptive manner.

DVAs are also essential for enabling multi-tenancy. Each DVA functions as an autonomous virtual network device with distinct performance characteristics and functional attributes.

Elastic Services Manager

The Elastic Services Manager is a virtual procurement and provisioning tool that streamlines deployment of DVAs. The ESM accelerates and simplifies provisioning of network services by automatically allocating and managing physical resources across DVAs. It eliminates management complexity and expedites deployment by automating upfront provisioning and configuration tasks and implementing dynamic, run-time changes. The ESM automatically configures the underlying physical x86 hardware and virtual resources, allocating Compute Units to DVAs based on administratively-defined policies.

The ESM helps service providers contain OPEX and improve time to market. It simplifies capacity planning and network engineering tasks and lets operators turn up customers, and define new services or grow and shrink service capacity quickly and easily – without involving the underlying server or VM infrastructure. Using the ESM, providers can efficiently provision DVAs and deliver consistent network services across diverse physical and virtual environments in a transparent manner.

RESTful APIs

Embrane offers comprehensive, easy-to-use RESTful APIs that enable straightforward integration with internally-developed or commercial service orchestration tools as well as other OSS and BSS applications. Complete APIs are provided for DVAs as well as the ESM.

Embrane has taken a radically different approach to developing application programming interfaces. Most layer 4-7 network appliances were originally conceived for the enterprise – where API requirements are minimal – and later re-purposed for hosted service or managed service applications. As a result, most network appliance vendors implemented APIs as an afterthought. They typically introduced APIs after-the-fact, exposing only a subset of features to programmers, often in the form of low-level, proprietary interfaces, or complex SNMP-based schemes.

Unlike conventional network appliances, the Embrane Virtual Network Services Platform was specifically created to enable network services under the cloud. Embrane recognizes cloud providers rely on a collection of third-party orchestration and provisioning tools to turn up customers and manage services. To that end, Embrane engineered into the platform a fully-functional, high-level, standards-based API. The API was designed into the product from day one, and is used internally for all management and control functions. As a result all DVA and ESM features and functions have been fully exposed in the API, and the interfaces have been fully exercised and validated during the normal course of product development and system test.

Distributed Virtual Appliances Redefine Networking in the Cloud

Embrane's distributed virtual appliance architecture redefines the way network services can be delivered by cloud providers. By decoupling network services from the underlying server and VM infrastructure, DVAs enable greater agility and higher scalability. Conceptually, one can think of a DVA as a virtual chassis that is instantiated across Compute Units. The DVA mimics the design principles of a physical appliance, yet has no actual physical dependencies. To the outside world, the DVA behaves like a physical device with distinctly-addressable network interfaces, management interfaces and application programming interfaces. Administrators can dynamically expand or reduce DVA capacity in minutes and in a transparent fashion, without reconfiguring hosts or VMs, re-hosting customers, or impacting management systems.

Figure 4 – Distributed Virtual Appliance Architecture

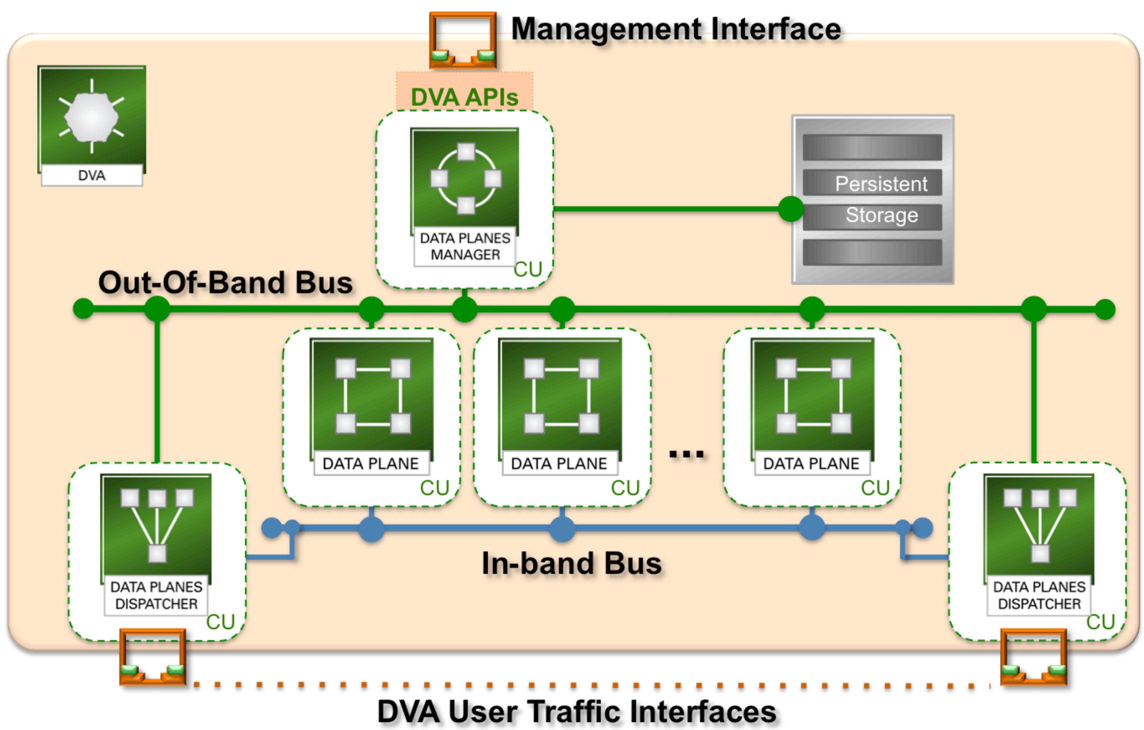

Distributed Virtual Appliance Architectural Elements

Architecturally, a DVA is similar to a traditional physical appliance. As shown in Figure 4, it includes an in-band bus for transporting user traffic and an out-of-band bus for transporting management and control data. However, in the case of the DVA, these buses are implemented as logical constructs that leverage the existing network infrastructure of the data center and not as physical connections within a chassis.

Compute Units are VMs that are allocated to the various tasks associated with overseeing and implementing network services. They interface with each other over the internal communication buses and with the outside world via externallyvisible network interfaces and management interfaces. The ESM automatically instantiates, configures and manages all internal DVA elements and communication paths, shielding administrators from the underlying implementation details.

Compute Units

A Compute Unit can take on one of several roles. Data Planes (DPs) are the core CUs responsible for carrying out layer 4-7 network services. They examine, process and forward packets based on administratively-defined policies. DPs can be added or removed from a DVA in a seamless fashion to address evolving capacity needs.

Data Planes Dispatchers (DPDs) implement the equivalent function of a network interface in a physical chassis and are responsible for directing inbound packets to DPs as well as forwarding outbound traffic outside of the DVA. A DPD advertises an external IP address to the outside world over a specific network interface. It uses affinity[2] rules to steer user traffic to a DP and ensure all subsequent packets associated with that flow are directed to the same DP. The affinity concept is instrumental for enabling scalability. It allows each DP to maintain a limited amount of state information, minimizing overall system synchronization requirements. DPDs employ a unique hashing technique that prevents them from having to maintain flow state and enables the dynamic insertion of DPs.

The Data Planes Manager (DPM) is the master of the DVA cluster. The DPM oversees the internal operation of the DVA and distributes policy and configuration information to the DPs and DPDs. The DPM advertises an external IP address to the outside world over the DVA's management interface. The DPM is the only CU that requires persistent storage. The storage is used to retain software images, configuration files, log files and other management-related data. Other CUs are diskless. As a result they can be provisioned rapidly – for both initial deployment and for failure recovery. And since they store no data they can be decommissioned in a secure and straightforward manner. DPs and DPDs boot over the network from the DPM.

Internal Communication Buses

A DVA includes two internal communication buses – an in-band bus for transporting user traffic and an out-of-band bus for carrying management and control data. User traffic is tunneled across the in-band bus to enable topology-independence, i.e. the CUs attached to the in-band bus can reside on the same physical host, on layer 2-adjacent hosts, or layer 3-adjacent hosts. Similarly, the out-of-band bus is a logical entity that is used to transport internal management and control traffic between the DPM, DPDs and DPs. The internal communication buses are not visible to the outside world.

External Interfaces

A DVA exposes several network interfaces to the outside world. End-customer traffic reaches the DVA via network interfaces terminated on the DPDs. Management traffic reaches the DVA via a dedicated management interface, terminated on the DPM.

---

[2] Affinity rules describe the relationships between traffic flows and DPs in a DVA. To limit the amount of state synchronization required between DPs, all packets belonging to the same flow are always sent to the same DP for processing.

Powering the Agile Network

Embrane is innovating the way network services are delivered, helping providers and enterprises meet the changing demands of the cloud. The Embrane Virtual Network Services Platform is the industry's first and only network services solution specifically designed with the cloud in mind. Embrane's distributed virtual appliance architecture combines the performance and scalability advantages of a hardware-based appliance with the flexibility and simplicity of a virtualized, software-based solution, enabling the delivery of extensible layer 4-7 on-demand network services with greater agility and higher performance.

Embrane Virtual Network Services Key Features and Advantages

- Easy-to-deploy, software-based implementation
- Hardware and hypervisor independent
- Open RESTful APIs
- Versatile, multi-service platform
- Standards-based solution
- Wide capacity range with elastic scale
- Dynamic, on-demand provisioning
- Multi-tenancy with service level guarantees

Forward-looking enterprises are pursuing cloud computing strategies to contain OPEX and CAPEX, improve business agility, and reduce operating risks. Embrane helps public cloud service providers and enterprises deploying on-premise private clouds deliver the critical on-demand networking capabilities that are essential for spurring cloud adoption.

To learn more about how Embrane can help your organization enable profitable on-demand network services visit us on the Web at www.embrane.com.

About Embrane

The networking landscape is changing rapidly and Embrane is at the forefront of the shift. Embrane's mission is to accelerate the adoption of cloud computing and IT-as-a-service by delivering the most dynamic platform for virtual network services. At Embrane, we are networking people solving one of networking's biggest problems. The network will finally be agile.

Embrane, Inc.
2350 Mission College Blvd.
Suite 703
Santa Clara, CA 95054
USA

www.embrane.com

embrane
Powering the Agile Network