

Cloud Computing Overlay for Unified Service Delivery: Delivering Infrastructure-as-a-Service

What You Will Learn

A top priority for many service providers today is entering the market for cloud computing services. A number of service providers are starting with “Infrastructure-as-a-Service” (IaaS) offerings that enable their customers to run workloads on the provider’s infrastructure. Such offerings have also been called “virtual infrastructure hosting” services.

In order to help service providers bring cloud computing services to market, as well as to improve delivery of traditional services in their portfolio, Cisco has developed the Unified Service Delivery solution. This solution combines a baseline architecture that is common to all services with a series of overlays that define how each specific service, such as IaaS, uses this baseline in concert with service-specific elements.

In this document, you will learn:

- How IaaS is delivered using this system of a baseline architecture and a service-specific overlay
- Which parts of the baseline architecture are required to deliver IaaS
- What service-specific infrastructure is described in the IaaS overlay
- How the Cisco® Unified Service Delivery solution provides distinct advantages in delivering IaaS

The Race to Cloud Computing Services

Cloud computing is the hottest topic in business services today. Providers must quickly roll out cloud computing services or risk missing the window of opportunity to define their place in this new market. But cloud computing is a broad term, and many providers do not have the applications or software development expertise to create Software-as-a-Service (SaaS) or Platform-as-a-Service offerings. Infrastructure-as-a-Service, on the other hand, is application-independent and provides an easy entry point for many service providers who are already adept at provisioning, managing, and scaling infrastructure-based services for their clients.

The challenge is that most infrastructure-based services, such as managed network services, hosting, and colocation services, are rooted in physical infrastructure. Cloud computing and IaaS changes the focus from the physical infrastructure to a new world of virtual infrastructure. No longer are services built around floorspace, servers, or ports, but rather around virtual machine (VM) images, application instances, and virtualized storage and networks.

Building and operating a virtual infrastructure can be a complicated undertaking for a service provider. Often, existing physical infrastructure cannot be easily repurposed to support a virtualized environment. Therefore, many providers are working to determine which new infrastructure components best support cloud computing services such as IaaS and other virtualized environments.

Further, the market for cloud computing services is far from mature, and industry experts expect continued rapid changes in the landscape. Therefore, decisions about infrastructure must stress flexibility of that infrastructure in order to accommodate whatever changes are on the horizon for cloud computing services. Those decisions must also reflect capabilities of the infrastructure to support non-cloud and non-virtual environments in cases where cloud computing and IaaS will coexist with traditionally delivered services.

Unified Service Delivery: Business Benefits

Cisco's Unified Service Delivery solution provides a single horizontal infrastructure that spans a provider's entire service portfolio. Doing this provides three main advantages to the business in general:

- Increases service and feature velocity
- Optimally utilizes capital and operating assets
- Helps ensure and secure the user experience

Specific to IaaS offerings, the Unified Service Delivery solution enables providers to:

- **Bring IaaS to market more quickly.** Virtualization-aware intelligence is built into the key elements of the Unified Service Delivery solution.
- **Reduce the cost of deploying IaaS and other services.** The Unified Service Delivery solution incorporates consolidation and scalability features that are unique to Cisco products.
- **Meet customers' service-level and security requirements in IaaS offers.** Quality of service (QoS), encryption, and secure partitioning technologies work in concert across the components of the Unified Service Delivery solution.

In addition, the Unified Service Delivery solution constructs IaaS offerings in a modular fashion on top of a baseline architecture that is not tied to any particular service. This had the added benefit of providing tremendous flexibility. Should the market for cloud computing services change, the underpinnings of the IaaS offer can be repurposed to support a substantially different service in response to shifting customer requirements.

Similarly, that same baseline infrastructure can be extended to other services beyond IaaS, even while it continues to be used as the foundation of an IaaS service. More traditional hosting or colocation services can be delivered over this infrastructure. Or non-cloud services such as hosted collaboration services can be delivered alongside the IaaS offer but using the same infrastructure. This reduces overall capital costs, increases the utilization of the infrastructure, and simplifies the environment operationally.

Cisco Unified Service Delivery for IaaS

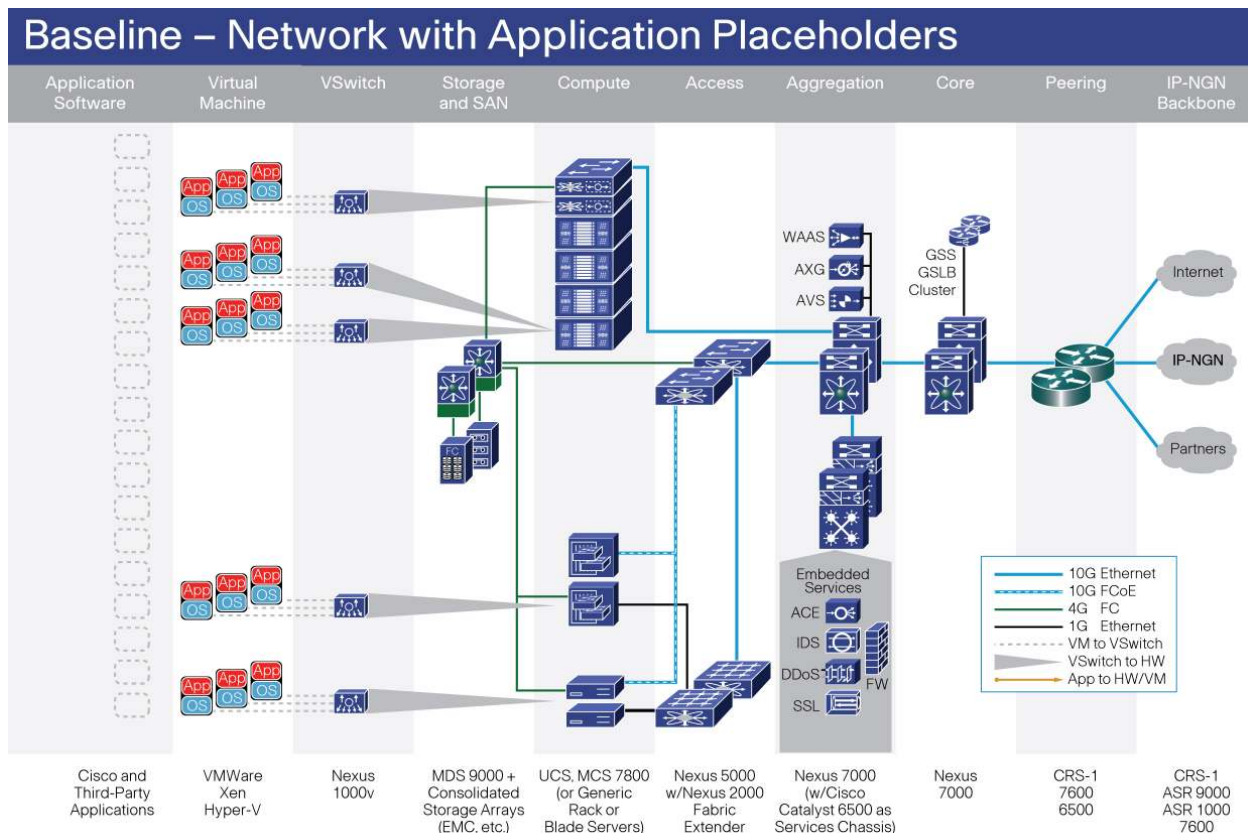
Overview of the Baseline Architecture

The baseline architecture for the Cisco Unified Service Delivery solution consists of network, storage, and compute resources deployed in a uniform fashion so that it is able to support a broad range of applications and services across the provider's portfolio. It uses leading-edge technologies from Cisco including:

- The Cisco Nexus™ Family of unified fabric switches for a single fabric of all traffic within the data center
- The Cisco CRS-1 Family of carrier-class routers for peering and interconnect functions
- The Cisco Unified Computing System to provide a compute-enabled, VM-ready network that is fully integrated with server virtualization solutions from partners such as VMware
- The Cisco MDS Family of storage networking equipment to facilitate storage consolidation in conjunction with partner solutions from companies such as EMC
- Cisco's Application Networking Services and security portfolios for virtualized, network-based services at Layer 4 and above

The components in the baseline architecture are depicted in Figure 1 as they would be deployed for a large-scale deployment that is very broad in scope, covering multiple services including IaaS, and supporting a heterogeneous multivendor environment in the compute tier. For smaller deployments, those with a more targeted focus, or more homogenous environments such as would be possible in greenfield scenarios, a number of the elements of the baseline architecture could be consolidated or eliminated.

Figure 1. Cisco Unified Service Delivery: Baseline Architecture



The Application Software Tier

On the far left of Figure 1 are placeholders for various applications or workloads that run in the provider's virtual environment. In most cases, the applications are supplied by subscribers as a preconfigured VM image. But applications also include the software infrastructure that the provider implements in order to enable the IaaS service. For instance, applications that provide the automatic resource allocation or implement subscriber self-service portals will be shown in this tier. Because this is the baseline architecture, Figure 1 shows "placeholders" as dashed boxes for where these applications will eventually appear when the overlay specific to IaaS is added on top of the baseline.

The Virtual Machine Tier

In a virtual server environment, multiple VMs, or logical servers, share the common resources of a physical server. The Unified Service Delivery solution relies on third-party technology from partners such as VMware's ESX to implement a virtual server environment. This technology, called the "hypervisor," is what enables the provider to abstract the virtual computing environment presented to subscribers from the physical servers that implement that environment. The hypervisor is also responsible for arbitrating access to the shared resources of the physical server.

With Unified Service Delivery, service providers can choose to deploy their VMs on the Cisco Unified Computing system, or on third-party servers. This is discussed more fully in the section on the Compute tier.

The vSwitch Tier

One portion of the shared resources of the physical server is the set of network interfaces that connects the physical server to the rest of the environment. There is not a one-to-one correspondence between physical network interfaces and logical network connections to the VMs. As a result, the hypervisor layer typically implements a virtual switching tier in the software of the hypervisor to multiplex connections from VMs to the physical network interfaces. In VMware, this is called the vSwitch.

This software-based virtual switching tier works like a very basic Layer 2 switch. But it is not managed with the rest of the network, so represents a blind spot for network administrators troubleshooting problems all the way back to the virtual machine.

Further, a service delivery environment almost by definition requires features to enforce security and service-level adherence. But a typical vSwitch is not an intelligent switch at all, and does not implement QoS or port security features.

To address these issues, Cisco has worked jointly with VMware to develop a replacement for the basic vSwitch that implements in software some of the essential network intelligence found in other leading-edge Cisco products. The Cisco Nexus 1000V Series Switch is a replacement for the standard VMware vSwitch available in vSphere 4. The Cisco Nexus 1000V Series incorporates technology unique to Cisco called VN-Link, which makes the network aware of server virtualization and thereby enables it to intelligently apply policy for configuration, quality, and security control not otherwise possible. While the standard vSwitch in VMware or another hypervisor can be used at this tier, the Cisco Nexus 1000V Series brings functionality that is expected of the physical network into the domain of the virtual environment, enabling IaaS providers to better meet the requirements of their enterprise customers.

Storage and SAN Tier

Moving left to right in the baseline architecture diagram (Figure 1), the storage and storage area network (SAN) tier is the first tier of physical resources depicted. This tier includes equipment designed to decouple storage from the compute tier and consolidate it for maximum efficiency. Also in this tier is the networking equipment required to connect the now decoupled storage capacity from the processing capacity in the compute tier.

This process of decoupling storage capacity from compute resources is particularly important in the virtualized environment of IaaS. In order for subscribers to be able to run their workloads on a VM that may be located on a physical server anywhere in the service provider's data center, it is essential that storage capacity is not constrained to storage devices (such as hard drives) that are directly connected to a particular server. If this were to be the case, subscribers could only execute their workloads on VMs running on physical servers that contained the storage devices where their data was located. That would severely limit the provider's ability to scale processing capacity on demand, to scale storage capacity beyond the limits of a few devices, to provide VM mobility to other facilities, and to provide other benefits that a truly virtualized environment delivers.

In the storage and SAN tier, the Cisco Unified Service Delivery solution incorporates both Cisco products as well as products from partners such as EMC. The Cisco MDS Family of Fibre Channel storage networking directors and switches enable IaaS providers to construct a robust SAN that permits flexible interconnection between the compute tier and storage in arrays provided by partners such as EMC. The Cisco MDS Family includes innovative features unique to Cisco and critical to developing a service delivery infrastructure for IaaS. Some of these features include VSANs, which allow a single set of physical storage networking equipment to implement multiple virtual SANs, and fabric-based storage applications such as storage media encryption for secure storage of data-at-rest and storage virtualization to further simplify operation and maintenance of the storage environment.

In many instances, a SAN environment will itself consist of a multi-tiered Fibre Channel environment. The Cisco Unified Service Delivery solution can incorporate such traditional SAN architectures, but a collapsed SAN is depicted in the illustrations to bring the focus onto the benefits of a unified fabric within the data center. This permits a single access tier to serve both the storage and application environments, and as a result, the diagrams show simple a SAN core with storage attached directly to the core. More detail about unified fabric is provided in the description of the access tier.

Compute Tier

A number of IaaS offerings provide computing capacity on demand, at variable scale, independently from the physical infrastructure required to deliver it. Ultimately, however, the provider must deploy real servers with real

processors, memory, and network interfaces, in order to deliver this capability. Those physical assets are what are included in the compute tier of the Cisco Unified Service Delivery solution.

The Cisco Unified Service Delivery solution incorporates recent innovations by Cisco in the compute space by including the Cisco Unified Computing System (UCS) in the overall solution. The Cisco UCS is a pre-integrated system that brings together advancements in networking, server virtualization software, memory, storage, and internal server architecture. It works transparently with the other elements of the Unified Service Delivery solution as well, including unified fabric for a single transport for all traffic within the data center, and the Cisco Nexus 1000V Series hypervisor-integrated virtual switch for extension of network intelligence into the virtual domain.

However, the Unified Service Delivery solution does not rely on the UCS. While the UCS makes integration and management within the compute tier significantly easier for providers, the Unified Service Delivery solution can be deployed in conjunction with existing server infrastructure. The Cisco Nexus 1000V Series software switch, Cisco Catalyst® Ethernet blade switches, and Cisco MDS Fibre Channel blade switches all enhance third-party servers in the compute tier, and can be integral parts of the Unified Service Delivery solution in heterogeneous environments.

Access Tier

At the access tier, a distributed switching infrastructure provides scale-out capabilities for the provider's IaaS offering. One of the complexities in delivering IaaS can be managing separate networks to accommodate both storage and application networking. The Cisco Unified Service Delivery solution addresses that problem by incorporating unified-fabric capabilities at the access tier. This permits both storage and application traffic to utilize a single transport while helping ensure that traditional performance characteristics of both environments are maintained.

Cisco implements unified fabric in the Cisco Nexus 5000 Series Switch, which is a key component of Unified Service Delivery. Unified fabric takes advantage of the increasing availability of 10-Gbps Ethernet interfaces on servers and the emergence of Fibre Channel over Ethernet (FCoE) as a means to transport storage traffic over Ethernet while preserving all of the performance and application characteristics of traditional Fibre Channel. As a result, a separate access tier for storage networking (the SAN) can be eliminated, collapsing two networks into a single unified fabric. The Cisco Nexus 5000 Series Switch implements a gateway function to permit transparent interconnection of the unified fabric with older SANs. This enables a provider to employ existing storage infrastructure, including storage arrays, and maintain established relationships with storage partners in an IaaS offering, even as native FCoE interfaces are currently in development.

The Cisco Nexus 5000 Series is a 10-Gbps switch, but to accommodate older LAN environments that continue to employ 1 Gbps for LAN traffic and a separate Fibre Channel or other storage environment, the Cisco Nexus 2000 Series Fabric Extender can be used to extend the LAN switching fabric of the Cisco Nexus 5000 Series to 1-Gbps ports housed in external fabric extenders. This provides a smooth migration path for existing provider LAN infrastructure being incorporated into the IaaS infrastructure. Of course, traditional Cisco switching technology in the Catalyst line can also be used at the access tier, as can traditional SAN technology as is found in the Cisco MDS Family of Fibre Channel switches.

Aggregation Tier

The aggregation tier in Unified Service Delivery provides a highly reliable, scalable mechanism for bringing together the traffic from the access tier and doing so in a manner that makes optimal use of physical resources. The aggregation tier is also the tier at which network-based virtual services, such as firewall services, load balancing, intrusion and anomaly detection, deep packet inspection, and other upper-layer services are typically inserted.

The cornerstone of the aggregation tier in Unified Service Delivery is tight integration between the Cisco Nexus 7000 Series modular switch, which provides industry-leading throughput in the data center, and the Catalyst 6500 Series Switch employed as a services chassis in which virtualized network services can be implemented. The Cisco Nexus

7000 Series delivers high-density 10-Gbps connectivity with up to 512 10-Gbps ports and over 15-Tbps throughput across the backplane, plus roadmaps to 40-Gbps and 100-Gbps connectivity, as well as modules supporting unified-fabric connectivity. With features such as these, the Cisco Nexus 7000 Series is ideally suited for this position in the network as server connectivity advances to 10 Gbps and the industry converges on a single data center transport standard.

Core Tier

In order to provide for the massive scalability projected as a requirement for some IaaS deployments, the Unified Service Delivery solution also includes a core tier within the data center, which is separate from the aggregation tier and also separate from the tier that peers the data center with the WAN. Naturally, in environments with lesser scale, this tier can be physically collapsed with either the aggregation tier, the peering function, or potentially both. But from a functional perspective, it is important to recognize that there are instances where a distinct set of characteristics defines the core of the service provider data center environment.

Within this tier, service separation is extended from the IP Next-Generation Network (IP NGN) into the data center using Virtual Device Contexts on the Cisco Nexus 7000 Series data center switch. This helps ensure that services can be isolated from one another at a very high level, so that current operational practices designed around line-of-business boundaries remain intact, while permitting consolidation of the underlying infrastructure. The Cisco Nexus 7000 Series includes high-availability features such as in-service software updates (ISSUs) and graceful process restartability – carrier-class features further enabling consolidation and massive scalability. Naturally, these capabilities extend into the aggregation tier implemented on the Cisco Nexus 7000 Series as well.

Peering Tier

On the boundary between the provider's data center and the WAN is the peering tier. Robust Layer 2 and Layer 3 feature sets are required here, as are a range of flexible LAN and WAN interface options, deterministic performance, and a range of platform options to support environments both large and small. As IaaS offerings may be deployed into environments already supporting existing services, it may be advantageous to provide complete physical and logical separation at this tier, effectively enabling two logical networks to be consolidated on a single platform – one network supporting the older environments and the other the future cloud computing services, as an example.

The Unified Service Delivery solution uses the Cisco CRS-1 router as the peering router because of its unmatched performance characteristics and feature set, which are ideally suited for the peering application. The Cisco CRS-1 Family has a proven track-record of service in demanding carrier core networks, and it has the same availability features as the Cisco Nexus 7000 Series, including ISSUs and graceful process restartability, making the two platforms very well matched in the overall solution.

Each platform is specifically tuned for the particular position it occupies in the network hierarchy. The Cisco CRS-1 router has WAN interfaces, protocols, and features designed to tie it with the rest of a provider's core network, or with other providers' networks or the Internet at large. The Cisco Nexus Family, on the other hand, is designed with features specific to the data center, such as unified fabric. By avoiding a "one-size-fits-all" approach, the Cisco CRS-1 is equipped with characteristics precisely aligned with the peering application – the right tool for the job.

IP NGN Backbone

While not a tier within the data center itself, this is nonetheless a critical component of the Unified Service Delivery solution. In this tier are networks including the provider's private backbone (if the provider is a network operator), partners' networks, and even the public Internet. Unified Service Delivery accommodates IaaS offerings from providers of all types – network operators or not. If the network to which the data center is attached happens to be a Cisco technology-powered IP NGN, there are additional benefits that an IaaS provider can add to the service by tightly coupling the data center with the IP NGN – either the provider's own or a partner's.

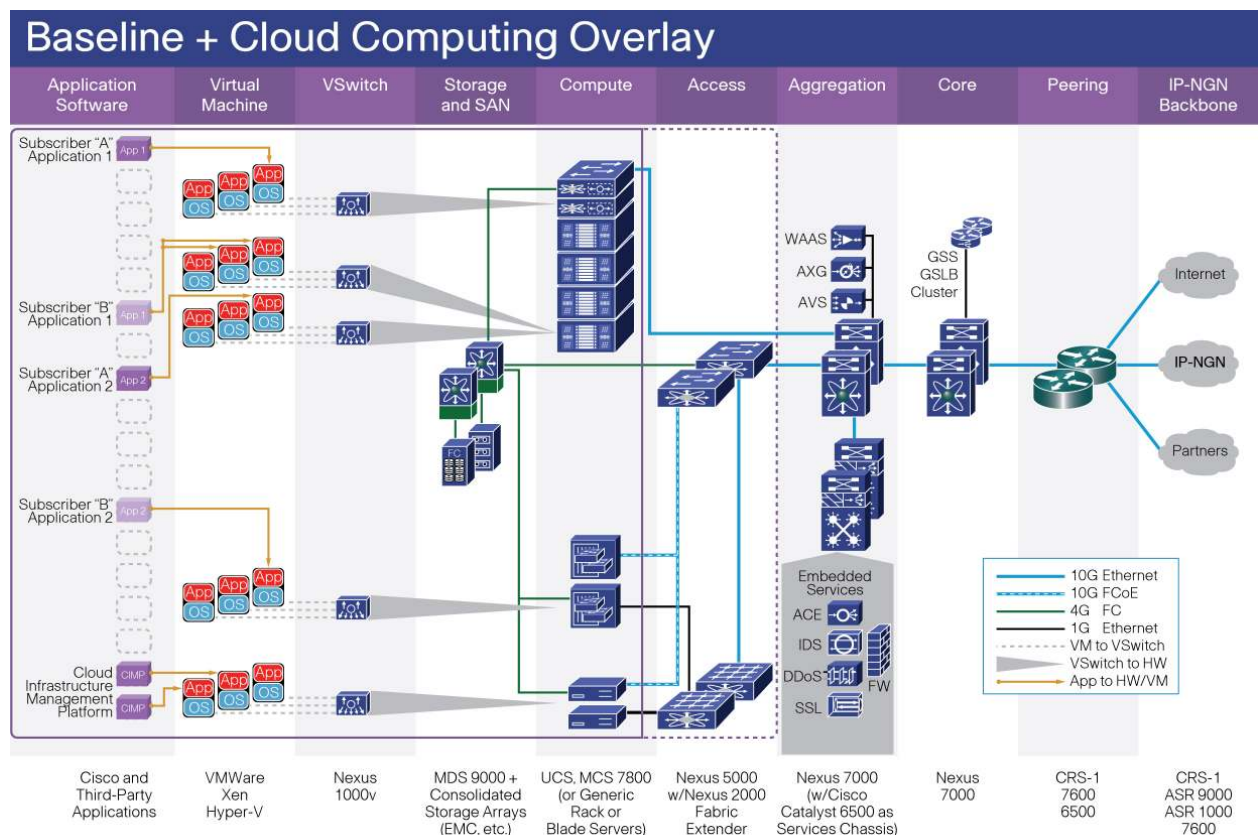
Some of the key elements of a Cisco IP NGN include the Cisco ASR Family of aggregation and access routers, as well as the workhorse of IP-based service delivery, the Cisco 7600 Series, and of course, the Cisco CRS-1 Family serving in core or edge applications. Features such as application acceleration, caching, network-embedded security, deep packet inspection, and service-quality monitoring are just a few of the examples of the added value that the Cisco IP NGN can bring when coupled with a Cisco data center in Unified Service Delivery.

Overview of the IaaS Service-Specific Overlay Architecture

In order to implement an IaaS offering, service providers can take advantage of many of the features of the baseline architecture, as described earlier. But there are certain specific elements that need to be taken into consideration when deploying IaaS that are not available in the generic baseline architecture, and generally are relevant only to the IaaS service.

Figure 2 depicts an IaaS offering deployed on top of the Unified Service Delivery baseline architecture.

Figure 2. IaaS Offering Overlay on Baseline Architecture



Cloud Infrastructure Management Platform

Perhaps the most critical element of an IaaS offering would be an infrastructure management platform. This is depicted in the lower right-hand corner of the drawing. This platform is what enables the provider to offer portions of the virtual environment to subscribers – on demand, at variable scale, and independent of the physical infrastructure – as if that environment were dedicated to the particular subscriber.

The infrastructure management platform is not a single piece of hardware or software, but rather a combination of offerings from Cisco, third parties, and the provider's own internal capabilities. It incorporates operational support systems, billing systems, self-provisioning tools, and element managers, all working in concert. The solid and dashed purple lines show the typical scope that a cloud infrastructure management platform might encompass.

With the Cisco Unified Service Delivery solution, common tools from Cisco and third parties that would be used to implement a cloud infrastructure management platform are tested in advance to help ensure that they interoperate in a way that most easily integrates into a typical provider's environment. This minimizes the effort required of the provider to perform such testing independently, and permits the provider to focus on the tasks that add value to the particular service they will bring to market. Common tools from Cisco include the Cisco Unified Computing Manager, Data Center Network Manager, Fabric Manager, Active Network Abstraction, and others. Third-party tools include vSphere from VMware, Control Center from EMC, etc., and would naturally vary depending on the specific provider's environment.

As illustrated by the mapping of applications that constitute the cloud infrastructure management platform onto virtual machines in the VM tier of the architecture, it is possible to run the management applications on the very same infrastructure that is used to deliver services to customers. This contributes to the overall efficiency and scalability of the environment, and removes unneeded complexity.

Subscriber Workloads and Applications

Beyond the cloud infrastructure management platform, no additional elements are needed to deliver an IaaS offering because, at its core, the offering is simply providing infrastructure on which subscribers can run their own workloads or applications.

Subscriber workloads are depicted in Figure 2 as additional applications in the application software tier running side-by-side in the environment. Of course, to subscribers, the environment appears as if it is dedicated to their application alone, and the security and QoS features of the Unified Service Delivery solution are what make that possible.

Other Services on a Common Baseline

One of the main benefits of the Cisco Unified Service Delivery solution is the ability to use the infrastructure as a common platform across an entire portfolio of services, and not just for a single service offering, such as IaaS. Not only does this help maximize utilization and minimize capital and operating costs, but it also can accelerate time to market for new services because infrastructure is already in place, eliminating the need for ground-up infrastructure builds with each new deployment. It also provides a mechanism for migrating existing services onto this new infrastructure so that those services, as they require infrastructure upgrades, can take advantage of the Cisco Unified Service Delivery solution as well.

Figure 3. A Common Baseline Architecture for all Service Overlays

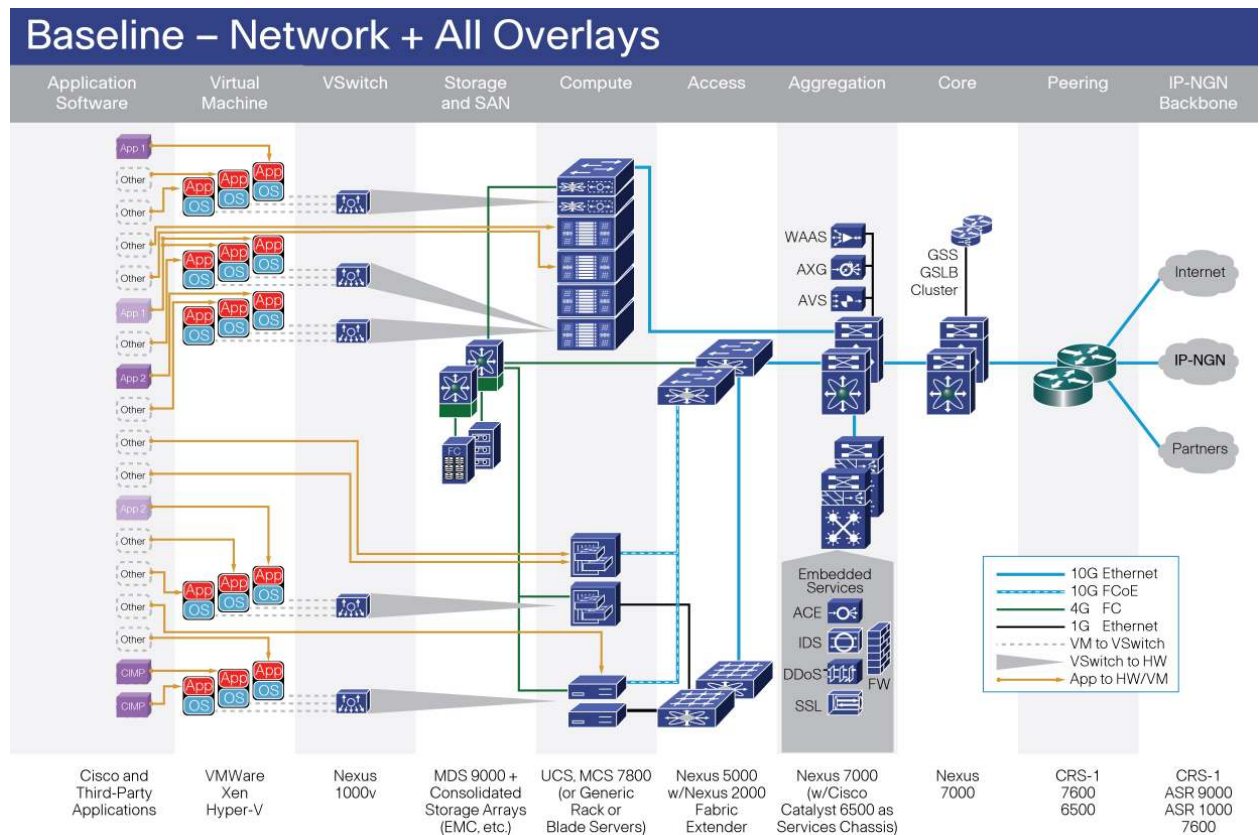


Figure 3 depicts how the Cisco Unified Service Delivery solution provides a common platform for all services across the entire service portfolio. Because the environment is completely virtualized, secured, and has mechanisms to help ensure QoS from application to subscriber, other services can run alongside an IaaS offering to make optimal use of the infrastructure. Examples of those services include video services for either consumer or business services, communications or collaboration services, or other cloud services such as Application-as-a-Service offerings.

Why Cisco?

Cisco has a product portfolio that spans the entire scope of service delivery infrastructure, from applications used for service delivery in the data center, through transport and intelligence in the IP NGN, out to endpoints for business and residential subscribers.

As a result, Cisco is in a unique position to assist service providers as they architect IaaS offerings that complement their particular service portfolio. Without this broad view, the approach to service delivery tends to focus on a few narrow areas, resulting in an environment that perpetuates silos and localized improvements while sacrificing the greater gains obtained from a fully optimized environment across the entire portfolio. Instead, Cisco takes a broad view of Unified Service Delivery, resulting in maximum value from the IP NGN and the Service Delivery Center (SDC) working in concert.

Cisco has proven this approach with our success in helping service providers transform separate, isolated networks into the IP NGN. This has spurred service innovation and the associated revenue growth for providers, improved profitability in service delivery, and raised the quality of experience for subscribers. Now Cisco is prepared to help its service provider customers repeat that success within the SDC, and tightly couple that environment with the IP NGN. This gives service providers maximum flexibility in the increasingly competitive market for traditional services, and positions them to capitalize on emerging trends such as cloud services.

For More Information

For more information, visit Cisco.com where the full range of Unified Service Delivery solution overlays are available. Also available are a set of consolidated links to the products and systems used in the Unified Service Delivery solution and customer case studies describing how Unified Service Delivery and the elements it contains have helped improve the economics and strategic value of their service delivery infrastructure.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)