



OPEN NETWORKING
FOUNDATION

OpenFlow™-Enabled Hybrid Cloud Services Connect Enterprise and Service Provider Data Centers

ONF Solution Brief
November 13, 2012



Table of Contents

2	OpenFlow-Enabled Software-Defined Networking
2	Executive Summary
3	Hybrid Cloud – the Next Generation
4	The High Cost of Overprovisioning
5	OpenFlow-Based SDN Solves for Mission-Critical Cloud Connectivity
8	Key Benefits
9	Conclusion
9	Contributors

CHALLENGES

ENTERPRISE

- Escalating IT operating expense versus declining per-user budgets
- Pressure on mission-critical data protection compliance, business continuance, and operational recovery times

CARRIER

- Weakening wireline revenues and profitability
- Growth of cloud service availability from over-the-top (OTT) providers, versus the dearth of new carrier service offerings

SOLUTIONS

- OpenFlow-based SDN hybrid cloud federation provides enterprises with self-provisioned transaction-aware data center inter-connect capacity with secure access to cloud provider's multi-tenant Infrastructure-as-a-Service

BENEFITS

ENTERPRISE

- Bandwidth-on-demand for timely, higher capacity inter-data center workload migration and processing, minus the burden of costly high-capacity fixed private line access

CARRIER

- Faster time-to-market and adoption of new, cost-optimized hybrid cloud services, and provisioning automation of responsive bandwidth-on-demand

OpenFlow-Enabled Software-Defined Networking

In a broad industry effort spearheaded by the Open Networking Foundation (ONF), Software-Defined Networking (SDN) is transforming outmoded network designs by decoupling the control and data forwarding planes, centralizing network intelligence and abstracting applications from the underlying network infrastructure using the OpenFlow standard.

As a result of this effort, enterprises and carriers will gain unprecedented network programmability, automation, and control, enabling highly scalable, flexible solutions that readily adapt to changing business needs.

Executive Summary

Traditional network architectures are not easily adapted to today's fast-changing and demanding enterprise and carrier business requirements.

While enterprise IT organizations face increasing pressure to reduce costs, often by way of infrastructure consolidation, carriers confronted with waning revenues need service offering innovation to attract greater numbers of enterprise customers seeking optimized services at reduced cost.

A decade of consolidation has demonstrated how compute and storage virtualization can dramatically reduce capital investment in enterprise IT, doing more with increasingly dense resources.

However, these changes were focused on data center server and storage virtualization, while the underpinning network architectures have stagnated with respect to both scalability and manageability. Networking innovation inertia has become increasingly problematic as the role of virtualization has transitioned from simple consolidation, to the aggregation of data center compute, storage and connectivity resources into a shared, but private cloud of services. These so-called private clouds have evolved to become agile providers of on-demand services to internal corporate customers.

Driven by business growth, mobility and the advent of Big Data, there is a new priority to intermittently augment private clouds with external resources known as provider-hosted, or public clouds. This in turn necessitates better connectivity between enterprise data centers and provider data centers—new and better ways to network private and public clouds.

The resulting 'hybrid' cloud infrastructure should easily provide transfer and sharing of data, or mobility of the applications working on that data, to support business and IT operations with scalability on demand.

Hybrid Cloud – the Next Generation

Opportunities have been created for OpenFlow-based SDN to help organizations build more deterministic, more scalable and more manageable virtual networks that extend beyond enterprise on-premises data centers or private clouds, to public IT resources that offer scale-on-demand to cost-conscious enterprise customers, while ensuring higher network efficiency to carriers seeking to improve their service profitability by provisioning more services, with fewer, better-optimized resources.

Hybrid cloud, where private meets public cloud, lets enterprises benefit from service provider economies of scale and scope, to lower their cost and improve application performance, resiliency and IT responsiveness.

At its simplest, this may come in the form of Storage-as-a-Service, or it may be as complex as automated compute and storage capacity expansion with intelligent placement of virtual machines into the cloud provider's geographically distributed data centers, resulting in distributed workflow processing among the public and private, or on-premise data centers.

To enable a range of hybrid cloud use-cases, there are four basic inter-data center, machine-to-machine functional building blocks:

- **Storage Migration**—making a copy of the data store in the enterprise data center and transferring the file(s) to the provider data center, either as backup, or as a precursor to active-active replication.
- **Active-Active Storage Replication**—keeping a data store in both the enterprise and provider data centers consistently synchronized, by writing through cache simultaneously to both locations.
- **Virtual Machine Migration**—transferring the CPU state and memory image of one or more applications running on virtual machines, from a server in the enterprise data center to one in the provider

SERVICE PROVIDER CATEGORIES

- Provide data center cloud services through their own data centers (e.g. Amazon, Terremark, Rackspace)
- Provide data center cloud services through their own networks (e.g. Verizon, ATT, NTT)
- Provide pure network services without data center cloud services (e.g. Level3, AboveNet)

data center, allowing dynamic server capacity expansion, or even fundamental business continuity through Active-Active data center or application failover.

- **Distributed Virtual Application**—communications between the virtual components that comprise a particular application instance such as between virtual machines or between a virtual machine and its associated data store, such as when web and application servers are placed in the cloud, but the database server is secured in the enterprise data center to ensure regulatory compliance.

These functions place demands on the hybrid cloud inter-data center network that must now support data transfers of 1 to 10 terabytes in as few as 1 to 10 hours, with roundtrip latencies of as few as 5 milliseconds, and a packet delivery ratio of 99.999% or better; essentially lossless.

As a result, meeting hybrid cloud bandwidth and latency demands, particularly those related to enterprise business continuity and disaster recovery, currently requires enterprises to lease high-capacity fixed links from network providers.

The High Cost of Overprovisioning

During the storage or virtual machine migration transactions at the beginning of a cloudburst into the provider cloud, bandwidth of 1 to 10 gigabits per second will generally be required. However, for the remainder of that IaaS instance life-cycle, much lower bandwidth, rarely exceeding 200 megabits per second, is required.

Currently, enterprises must contract for over-provisioned fixed capacity to meet the multi-gigabit peaks, which results in costly, underutilized capacity during sustained quiescent periods. Conversely, if they contract for under-provisioned capacity to meet lower cost targets, they risk protracted cloudburst operations that may even require physical shipment to get a data store into the cloud, and an inability to migrate live virtual machines without application interruption.

When forced into this no-win decision, IT can opt to over-provision to avert internal customer dissatisfaction or outright breach of SLAs with potential external customer loss—and this eliminates the hybrid cloud solution for the many enterprises that simply can't afford the high cost of overprovisioning.

Carriers have continued to limit their offers to fixed-capacity services because the manually managed, detailed configuration of traditional circuit-switched optical environments coupled with the lack of comprehensive connection admission control necessary to dynamically and reliably adjust committed information rate for virtual circuit switched Ethernet environments, has made remote provisioning anything but dynamic, increasing the complexity and cost of providing a flexibly-priced data connection service, and introducing delay between the user request and the delivered service.

While the traffic on the leased connection may be well below maximum capacity, the service provider has no automated mechanism to orchestrate the network to leverage available bandwidth and offer it to other customers.

However, with the option to automatically provision, adjust and de-provision services, carriers could grow their revenues by coupling that capability with a pricing structure to meet enterprise requirements with greater diversity. Correspondingly, enterprises could more easily adopt a hybrid model where dynamic network capacity was automatically or self-provisioned on-demand, and priced accordingly.

DEMYSTIFYING FABRICS

Fabrics are topologies often illustrated as a mesh due to the high density of connection paths required to ensure a set of connectivity characteristics between nodes. These characteristics include **deterministic** – for pre-determined path selection prior to transmission; **multipath** – where alternate paths are available; **self-healing** – where high failure rate paths drop to continuous test mode, and are restored automatically as the failure rate falls; **lossless** – where frame arrival is guaranteed; **low latency** – where transmission incurs the least delay; and so forth.

Originating with the crossbar switch fabric, Fibre Channel popularized the fabric moniker by meshing crossbar switches to provide the above characteristics to storage networks. Ethernet fabrics are becoming more common in the data center since DCBX standards (ETS/ PFC) were adopted in Ethernet switching silicon, however classic Ethernet cannot assure the above behaviors.

OpenFlow-Based SDN Solves for Mission-Critical Cloud Connectivity

Software-Defined Networking with OpenFlow provides the framework and tools to enable dynamic enterprise-provider data center inter-connect capacity that matches, and is directly driven by, the cost-benefit ratio of hybrid cloud service fulfillment. At a high level, the key solution elements shown in Figure 1 are:

- OpenFlow-enabled cloud backbone edge nodes that connect to the enterprise and cloud provider data center fabrics, and OpenFlow-enabled cloud backbone aggregation and core nodes which efficiently switch traffic between those edge nodes
- an OpenFlow-based SDN controller to configure the flow forwarding tables in those cloud backbone nodes, and supporting a network virtualization application (aka WAN hypervisor) to allocate the shared network resources among those nodes
- Hybrid cloud operations and orchestration software to manage the enterprise and provider data center federation, inter-data center workflow, and compute/storage and inter-data center network resource alignment




FIGURE 1
 The enterprise private cloud in data center A is connected to two regional data centers in a service provider’s public cloud, forming an enterprise-service provider hybrid cloud. Note that the data center networks are high-availability ‘fabrics’ that provide multipath, deterministic, guaranteed latency connections required by storage or compute clusters. The OpenFlow controller provides fabric ‘continuity’ across the regional optical network, ensuring all deterministic networking characteristics are maintained within the hybrid cloud to satisfy both the synchronous storage connections (Fibre Channel or 10GbE) or the <10 Gbps Ethernet Virtual Private Line connections used for asynchronous storage.

For an enterprise that signs up for hybrid cloud service, in the course of their IT operations there will be events that precipitate the advantageous movement of data store, virtual machine, inter-machine messaging and/or storage update traffic between physical servers and storage devices in their data center and those in the cloud provider’s data center (refer back to the hybrid cloud functional building blocks above).

This workload movement can be triggered explicitly through a user portal or automatically from IT performance monitoring and workload balancing software.

Both approaches initiate transactions by plugging into the hybrid cloud operations and orchestration software that resides in the cloud provider’s operations center and communicates with an associated client module in the enterprise (either provider-supplied or based on Open Data Center Alliance open standards).

THE STORAGE MIGRATION USE CASE

Application mobility has become a staple of cloud orchestration. However, as mobility has introduced networking complexity, the challenge of maintaining the tight coupling of distributed applications with the large volumes of distributed data they create has become a new and insistent priority.

Traditional network architectures are not well adapted to facilitate intermittent relocation of terabytes of application data, slowing or even preventing data migration over distance, and making the advantages of orchestration unattainable for large applications with unwieldy data sets. As a result, many large enterprise applications have remained immobile and inefficient.

Software-defined networks based on OpenFlow are essential to facilitate efficient application and data orchestration, and to extend the range and flexibility of private clouds beyond the confinement of a single data center. This simplifies enterprise connectivity between their remote data centers, letting them more easily form hybrid clouds with service provider infrastructure.

In an OpenFlow-enabled software-defined network, hybrid cloud services can orchestrate the migration of applications and data in partnership with a virtualized network, leveraging available storage replication or mirroring services. In this way, the network can be automatically programmed to provide the services required for the storage replication or migration service to work in conjunction with—and remain tightly coupled to—the application migration.

An OpenFlow-enabled network brings intelligence to connectivity virtualization, so that when physical resources necessary to complete data migration per the expected service level are not immediately available, the OpenFlow controller can suggest, schedule or automate resource availability to meet service levels based on pre-defined policies. The hybrid cloud orchestrator can then initiate the migration service when the OpenFlow controller announces resource availability.

Through orchestration software and domain hypervisors, virtual resource capacity is allocated on specific physical devices in the provider data center.

The network addresses of the source and destination host devices and the virtual compute and storage resource instances can be used to uniquely identify the inter-data center flow to transfer the workload.

The nature, magnitude and time constraint for completing the workload will dictate the inter-data center bandwidth and QoS required for the flow.

SDN enables the network to be dynamically configured to identify and accommodate this composite flow specification. The orchestration software communicates with the network virtualization application to request or negotiate for the bandwidth required between the two data centers to handle the uniquely identified flow.

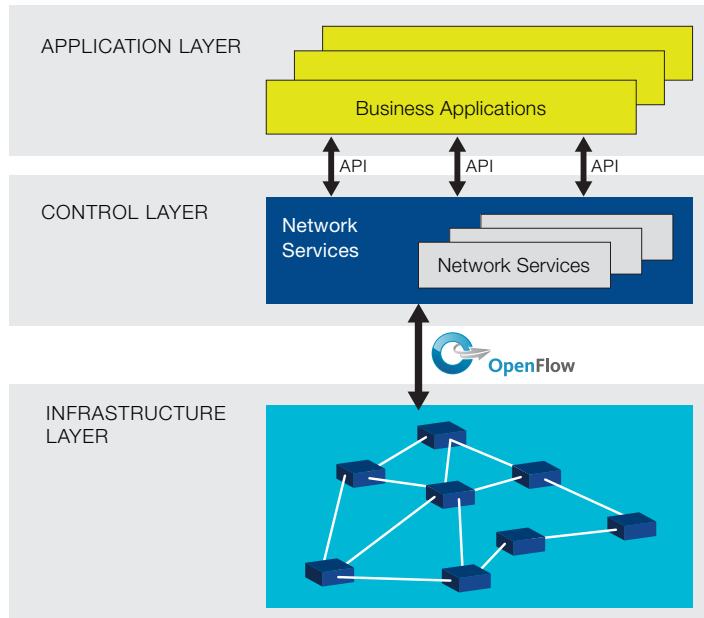


FIGURE 2
Software-Defined Network Architecture

Based on its global view of resources, if the requested resources are available and if the request adheres to the policy administered for this customer, then the underlying OpenFlow controller will use OpenFlow protocol messages to configure the fabric-facing ports of the cloud backbone edge node to admit and police the flow and schedule the associated packets in accordance with the flow specification to ensure the negotiated bandwidth and QoS.

This may be replicated at the cloud backbone aggregation and core nodes, or the traffic may be encapsulated and forwarded by the cloud backbone edge node into a preconfigured tunnel for tunnel switching by the aggregation and core nodes.

APIs enable 'northbound' messaging between the OpenFlow controller and the Cloud Orchestration, as shown in Figure 1, so that once the flow has been provisioned, notification is sent back to the orchestration software to initiate the inter-data center workload movement. When the transaction is complete the orchestration software will request the connection to be torn down and OpenFlow will remove that flow table entry from the cloud backbone nodes and the network virtualization application will provide event messages or statistics to the billing system.

Throughout the operation the network controllers use information about the current network usage to load balance links in the service provider network, ensuring that all data center tenants are receiving the service characteristics they require and that all network links are being used as efficiently as possible.

The efficient use of available link bandwidth allows service providers to maximize the bandwidth available for purchase by customers, improving their service offering while ensuring all SLAs are honored.¹

1. Optical transport networks are currently controlled using extensions to OpenFlow (shown as OTN/λ in Figure 1) that are not included in the OpenFlow 1.0 release, but are in consideration for a future release by the Extensibility Working Group, tasked to maintain the core of the OpenFlow switch specification.

Key Benefits

There are general advantages to be realized by organizations that adopt OpenFlow-enabled software-defined networking as the connectivity foundation for private and hybrid cloud connectivity.

A logically centralized SDN control plane will provide a comprehensive view of data center and cloud resources and access network availability and this will ensure cloud-bursts are directed to adequately resourced data centers, on links providing sufficient bandwidth and service levels.

Openness Drives Enterprise Cost Optimization

OpenFlow-enabled SDNs will facilitate multi-vendor networks between enterprise and service provider data centers, helping enterprise customers to

- choose best-in-class vendors, while avoiding vendor lock-in
- choose from a wider variety of access technologies (e.g. DWDM, DSL, HFC, LTE, PON, etc.)

EXTENDING THE OPEN NETWORK

OpenFlow standardization has focused to date on the packet infrastructure. Other layer networks may be supported in an SDN environment—i.e., one featuring a centralized control layer—through closed-system application of existing technologies.

However, continued evolution of the OpenFlow standard may extend open interface support to those additional technology and infrastructure layers which are critical to the WAN. Such developments would extend the utility and benefits of OpenFlow-based SDN to the full suite of layer technologies and services associated with end-to-end cloud interconnect services.

- access dynamic bandwidth for agile, timely inter-data center workload migration and processing
- ease or eliminate the burden of underutilized, costly high-capacity fixed private line leases

Openness Drives New Service Creation

OpenFlow-enabled SDNs will help carriers and service providers to free up and market underutilized bandwidth with differentiated services. OpenFlow-enablement ultimately will result in

- faster time-to-market for new services with an accelerated adoption cycle and attach rate
- more competitive positioning through SDN-enabled bandwidth-on-demand services and more dynamic response to infrastructure and application demands
- provisioning automation and intelligence driven by cloud service orchestration logic
- improved customer retention through value-added hybrid cloud services

Conclusion

The promise of OpenFlow-enabled Software-Defined Networking is that it will allow innovation to supersede the stagnation of networking technology, and this in turn will transform IT, making it more resource-efficient and more responsive to the needs of the enterprise and service provider alike.

Contributors

Mitch Auster, Editor

Nabil Damouny

John Harcourt

Open Networking Foundation / www.opennetworking.org

The Open Networking Foundation is a nonprofit organization founded in 2011, whose goal is to accelerate the adoption of open SDN. ONF emphasizes the interests of end-users throughout the Data Center, Enterprise, and Carrier network environments.

Open Networking Foundation, the ONF symbol, and OpenFlow are registered trademarks of the Open Networking Foundation, in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.