



Unified Network Fabrics: A Veritable Inflection Point

Application delivery has benefited from increases in efficiency, agility, and scale by using abstracted, scalable designs for computing in a variety of ways over the past five plus years. Advances in multitenant cloud and web scale applications have driven the gains to the benefit of users and IT service providers in many environments. In parallel there has been strong interest in bringing the benefits of abstraction, simplification, and integration with application platforms to network infrastructures delivering services. The vision of highly scalable, flexible, programmatically accessible networks has been a driving force in early stage designs in software defined networking over the past several years.

Progress, however, has only been achieved in relatively confined pockets of networking. In data center infrastructures, for example, gains have been made using more open hardware and software in switching platforms, producing lower costs of entry and new possibilities for software-based enhancements in the operation of the underlay physical network. In parallel, developments have brought agile and automated overlay virtual networks into play in support of multitenant clouds and hybrid physical and virtual infrastructures.

But in these developments, offerings have been constrained to addressing relatively narrow slices of the overall communications problem. They may support virtual machines in a cloud service environment without offering visibility into the physical underlay on which they depend and creating another operational island that needs its own specialty support. Or they may supply open networking software in a network OS in conjunction with merchant silicon-based hardware switches but may be oblivious from an operations point of view about the behavior of the overlay virtual network on which cloud computing resources depend to do their work.

KEY FINDINGS

Unified physical + virtual fabrics bring a new network model to virtualized, software-driven data centers.

Big Switch Networks' Big Cloud Fabric is the market's first unified P + V fabric based on open SDN principles at each level of implementation.

The 5 year total cost of operation (TCO) of BCF in a 16 rack, 640 server VMware cloud computing POD is 40% of the TCO of a comparable configuration based on present mode of operation (PMO).

BCF is 29.6% of the 5 year TCO of a PMO configuration in an 8 rack, 320 server OpenStack cloud computing POD.

Installing and activating the network for the cloud configurations is 16x faster using BCF's Zero Touch operations than in the node-by-node approach of the PMO.

Diagnosing path integrity between virtual machines is 15-18x faster using BCF's automated path trace than in the node by node query approach used in the PMO.

Realizing top-line benefits of feature innovations is on average 4x faster in the BCF implementation than in the PMO.

The result, in short, is a series of isolated gains inching toward a more unified solution but remaining in fact segregated and suboptimized against the ultimate goals of the industry.

Because of this fragmented progress, conditions are ripe for delivery of a more complete solution to the data center networking problem—one that is agnostic to whether a networking function is realized as a virtual networking node or a physical networking device for transporting application traffic. Delivering a unified fabric that closes the gap between overlay and underlay systems, presents a continuous framework for delivering services on an application or per-tenant basis, and erases the need to dwell on lower level components of network operations is an idea whose time has arrived. Making such a fabric available would represent a quintessential inflection point in the evolution of networking that would be a challenge to reverse once achieved.

Big Cloud Fabric 3.0: Opening the Door to Simpler, More Efficient Operations

Big Switch Networks has harnessed these observations as fuel for its vision of a unified network fabric to support agile, efficient delivery of services in heavily virtualized and software-driven data centers. It has grounded the development of its Big Cloud Fabric (BCF) on principles that eliminate the barriers present in prior implementations and enable the efficiency and productivity in operation that implementers have been searching for. With the introduction of BCF 3.0 Big Switch is moving further still along this path.

BCF uses principles of abstraction and logical centralization that make realizing a unified network fabric for application services possible. Even though network infrastructures themselves are distributed (they run in a variety of nodes) it is simpler to control them if key functions they are doing can be seen and understood in a consistent manner from a logically centralized control point. This notion has been useful in developing large scale network protocols, modular chassis nodes, and many kinds of management applications in the past. It is employed in the BCF model by integrating key logical controls in the BCF Controller. In this model the fabric BCF is implementing can be seen as ‘one big switch’ in the logical sense. Similarly, abstraction of the network elements being operated on allows for simple, efficient steps to be taken in accomplishing a task. This allows for user groups and applications to be seen by the operator as just that (as opposed to, say, groups of VLAN IDs and link aggregation parameters). Abstracting terms of reference and the data models supporting them in the system allows for many efficiencies to be achieved.

These two principles allow a perspective to be adopted in which it doesn’t matter whether a network element is deployed as a physical or a logical node to be operated on efficiently. A unified view of how topologies are constructed and how policies are enforced is made possible and spans the physical and virtual domains. As part of making this possible BCF 3.0 introduces an overlay virtual version of Switch Light network software that works in tandem with BCF physical nodes, and managed from the logically centralized, resilient BCF Controller that coordinates the fabric overall.

In parallel BCF employs two dimensions of open architecture that also enable its advantages. First is its use of merchant silicon and open hardware designs that allow it to concentrate on software agility and content on top of open network underlays, and benefit from the cost advantages being brought to market by the merchant silicon ecosystem setting new benchmarks for networking infrastructure costs in many domains. Robust nodes can be delivered in this model at dramatically lower costs than has been achievable in proprietary, single supplier model of prior solutions.

The second dimension of open architecture leveraged extensively in BCF is open modular software at key points in the system's design that accelerates feature velocity as well as integration of the fabric into environments such as Hadoop/big data applications, multi-tenant clouds, virtual desktop infrastructures and high performance storage. By using APIs internally that allow for efficient development of 'local' functions, along with well documented and widely adopted APIs for things like 'northbound platform' integration (cloud and application management systems) and user-defined enhancements (in-house optimizations specific to a customer's business) BCF provides the agility and flexibility that is needed for networks to support the range of applications users want .. and do it with the simplicity and cost-effectiveness they also need. This broad embrace of software agility contributes heavily to achieving the goals of 'Zero Touch Operation' in unified virtual and physical infrastructures that Big Switch is pursuing.

Because delivery of unified P + V solutions is so new in the market overall, Big Switch determined it would be an excellent time to evaluate the difference in economic results that can be achieved when using the BCF architecture in data center deployments compared to prior modes of implementation, and to leverage the tools of analysis that have been productively used in architecture transitions of this type in the past.

Private Cloud Implementations: A Compelling Use Case to Analyze

Use cases that benefit from open SDN fabrics range from multitenant clouds to big data analytics platforms, high performance storage, and virtual desktop hosting. The case with which Big Switch has been engaging most frequently with its customers and prospects is the private IT cloud. Whether a purely private, single cloud computing POD or a larger deployment involving multiple PODs or sites, it is the deployment on which the largest number of IT teams has been focusing. It is also an environment to which BCF is well suited. For this reason Big Switch decided to focus on the private IT cloud as the target for in-depth analysis of benefits achievable using BCF compared to the prevailing mode of data center network deployment.

The two cloud computing frameworks most frequently considered for private IT clouds are VMware's vCenter-based solutions and various distributions of the OpenStack cloud computing software (from firms such as Mirantis and Red Hat). Because of this, one configuration using each of these cloud management software environments was analyzed in a Big Cloud Fabric (BCF) and a present mode of operation (PMO) deployment.

We used a single data center POD as the basis of the analysis to keep comparisons clearly scoped and understandable. The benefits achievable in a single POD can largely be generalized and multiplied in a near-linear fashion when extending deployment to multiple PODs or sites.

VMware Private Cloud Deployment

The VMware POD in our analysis comprises 16 racks, each one 42 rack units (RUs) in height, each housing 40 1 RU servers. The POD is thus a 640 server POD. This is a common size for a fully configured data center POD, though sizes vary by site and organization.

Both the BCF and the PMO solutions are deployed in the POD with two 1 RU leaf (or top of rack) switches per rack, allowing each server a redundant connection to the underlying network. Each leaf switch contains 48 10GbE server-facing ports and 6 40 GbE spine (or core) facing ports. Each leaf switch is connected to multiple spines to ensure availability and performance. Spines are connected only to

leaves, performing the function of an interconnect fabric as one would find in a conventional networking node. This configuration can be viewed as a leaf-spine CLOS, conveying the notion of a flat, high-performance fabric for the cloud. Connections to other domains (such as data center core or wide-area domains) is performed by leaves in this configuration assigned to connecting at those boundaries.

Additional details on the BCF and the PMO solutions in the VMware POD are provided in the section that details their economic differences.

OpenStack Private Cloud Deployment

The OpenStack POD in our analysis comprises eight equipment racks with 40 1 RU servers in each for a total of 320 servers (or similar devices) in the OpenStack cloud. This POD is smaller than the VMware POD not because of technical or architectural considerations related to BCF, PMO or OpenStack solutions. It is this size because that is the scale at which Big Switch has qualified BCF 3.0 with OpenStack. Going forward installations larger than eight racks and 320 servers with OpenStack are expected to be qualified by Big Switch for BCF.

Additional details on the deployment on the BCF and the PMO solutions in the OpenStack POD are provided in the section that details their economic differences.

Simply More Efficient

As context for thinking about the analysis, we describe certain architectural differences between the implementation of BCF and networking solutions predominant in the data center networking environment that have been in use during the past decade and longer. These architectural differences are central to understanding the differences in functional and economic performance that are achievable in either case.

BCF is based on an enhanced model that delivers an order of magnitude improvement in results compared to the mode of operation dominant in data center networks to date. BCF is well summarized as an open SDN fabric, but it is the elements of the architecture that enable its specific benefits.

BCF is modeled as a single unified logical communications fabric, one big switch, regardless of whether functions are performed in a virtual or a physical network element. This allows for consistent process to be applied to physical and virtual components in a unified P+V design. Most network solutions have mastered only a physical or only a virtual function and have not yet achieved the transparency of unified P+V. Dominant PMO solutions entail different controls and operations for the virtual and the physical domains, increasing overhead and cost and diluting overall results.

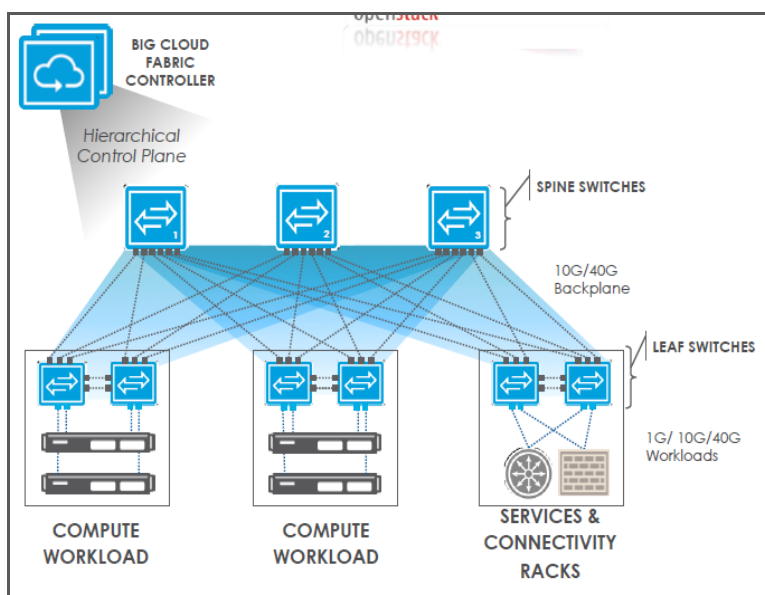


Figure 1. Single Logical Fabric Design of BCF

Big Switch has also embraced open hardware architectures leveraging merchant silicon and white box designs as part of its foundations. Following this path allows it to ride the cost curves being demonstrated now in open networking hardware, which have been approaching Moore’s Law refresh cycles in the past several years compared to the five-year plus refresh cycles of preceding proprietary chips. Keeping its focus on robust software functionality that can be tuned to the design of underlying merchant silicon designs gives BCF a sustainable capital expense (capex) edge as well as flexibility to adapt its software to different underlying platforms.

In a similar vein BCF is built using open software designs pervasively. Internally and externally modules are built on APIs that emphasize flexibility in adding functions for BCF on its own (adding a new forwarding function, for example) as well as for integration with external platforms such as cloud management systems such as vCenter and OpenStack.

BCF also leverages a logically centralized control model that brings significant simplification to a number of operational tasks compared to doing work on a node-by-node, function-by-function basis across a distributed network configuration. Combined with its extensive work on abstracting the description of elements being worked on and functions being performed, BCF is able to achieve a degree of automation in the execution of deployment and operational tasks that is simply not achievable in the more fragmented, less centralized and less abstract designs of PMO solutions.

Framework of Analysis

We concentrated primarily on total costs of operation (TCO) over a five year period when using the BCF or the PMO solution in the VMware and OpenStack PODs. We also included as a lower priority an example of incremental benefits in other areas of operation a company may achieve based on the ease or speed of innovation possible in the two alternatives, and the benefits that could accrue in revenue-generating, market-facing activities. This metric is focused on the value of innovations in the alternatives cases and produces a view of incremental cash flows achievable between them (this is the ‘net’ of costs to implement versus benefits achieved - it is a form of accelerated time to results).

Cumulative total cost of operation (cumulative TCO) includes both capital (purchase) costs and operational, workflow-based costs (operation expense) related to the alternatives over a defined period. In this analysis we compared cumulative TCO of the alternatives over five years. Additional aspects of the metrics we analyzed are described in other sections.

Capital costs include the cost of acquiring the hardware switches (leafs and spines) for the configurations, using current list prices of both BCF and comparable PMO devices. Capex also includes the cost of software licenses for the configurations. While methods of obtaining software vary and are likely to evolve, we chose this method to be clear on the relative costs of software in the different designs and to be consistent across the alternatives. In each case the cost of acquiring either a five-year or a perpetual license for the particular software package was used, these fitting the five-year framework best. Finally, we included the cost of hardware and software support in capex for this analysis. Sometimes support costs are considered operation expense (opex) because they can be incurred on an annual or other recurring basis. However, we chose this alignment so we could focus on the costs of acquiring the platforms in capex, and the opex analysis could focus on the IT staff and service delivery efficiencies that apply to the alternatives.

Opex in our analysis includes the costs of planning for, installing, configuring, operating, and expanding the alternative solutions (BCF and PMO). This is often viewed as a life-cycle or workflow-based approach to measuring operational costs and is focused on the relative efficiency and productivity IT and other teams achieve when using one alternative or the other. We excluded costs such as power, space, and cooling overhead in this analysis because the differences between the alternatives are minimal on those metrics, and the overwhelming basis of the advantages of the BCF design over PMO is in the capex, opex, and innovation-driven metrics.

BCF and PMO in a VMware-Based Private Cloud

Capital expenses for supporting the 16 rack VMware POD are summarized in Figures 2 and 3. At the top level **the BCF configuration costs 52.8 percent of the PMO configuration for the five-year analysis, all capital costs accounted.** A substantial portion of the difference can be attributed to the difference in cost of hardware between the merchant silicon of BCF switches and the custom silicon of the PMO. Costs for leaf and spine switches in each chart are indicators of this. Of interest is the relatively high cost of support in the PMO versus the BCF case. On one level this is because it is a disaggregated cost per device (PMO case), and one could argue that the BCF design channels a significant amount of support through the centralized architecture of the controller where the majority of the control functionality is homed. An equally useful observation is that BCF places higher value on the role of software in the fabric overall and allows for the software that defines its functionality to command a higher percentage of its overall price while streamlining the delivery of support. It helps amplify the value of observing that a solution with automation in its overall operation is also produced at half the capital cost of entry to the customer.




Figure 2. Capex of BCF in the VMware POD




Figure 3. Capex of PMO in the VMware POD

On top of the capex advantage, using BCF in the VMware cloud also brings definitive opex benefits versus the PMO approach. These benefits range from simplifying planning to using dramatically simplified software upgrade procedures to supporting more seamless diagnostics because of the tight integration of the underlay and the overlay configurations. We highlight the description of one of the benefit areas, focused on dramatic simplification of initial installation and activation of the network for the POD. Similar benefits accrue in initial installation in the OpenStack case as well, and we reserve another opex example that also benefits both POD scenarios for description in the OpenStack section to keep the discussion consumable and balanced.

Opex Simplification in the VMware POD with BCF: Automated Fabric and Auto-LAG Formation in the P+ V Fabric

Instantiating a Big Cloud Fabric and enabling a PMO network for a private cloud using VMware are two radically different processes. The overall comparison could be summarized by zero touch versus heavy touch. BCF uses the abstraction of a single logical fabric coordinated by a central controller along with simple references to components of the fabric described by their role (leaf or spine, for example) and a few intuitively defined relationships (for example, graphic alignment of nodes in a topology GUI including links, whose details can be automatically configured based on their location and role). BCF

takes these higher level constructs and extensively automates its configuration and activation processes to remove handling minutia of individual interface configurations from the engineer's tasks. Details of deployed configuration (port configuration details, for example) can be seen by an engineer who wants to see them (ensuring visibility of components to any level if the deployment team wants). But the default mode of operation is to have them simply happen as compared to enabling the configuration manually, one line at a time.

By contrast when installing the PMO leaf and spine configuration each switch in the POD needs to be individually configured one switch at a time to define its interfaces and control plane operations appropriately. **In the 38 switch configuration of the VMware POD this process is on the order of 16 x longer for configuration of a working PMO network underlay when compared to zero-touch controller-driven approach of BCF.**

Beyond the initial installation of the network itself, an important aspect of cloud deployments is the efficiency of integration with the cloud system's overlay virtual network and related resources (virtual machines, application segments, etc.). The goal most teams strive for is as seamless an operational fabric as can be achieved, blending underlay and overlay efficiently. Using logical abstraction, centralization, and zero-touch automation principles, BCF constructs a tight integration between itself and the VMware cloud to simplify configuration and increase operational efficiencies.

A view of how this is achieved in the VMware case is shown in the Figure 4.

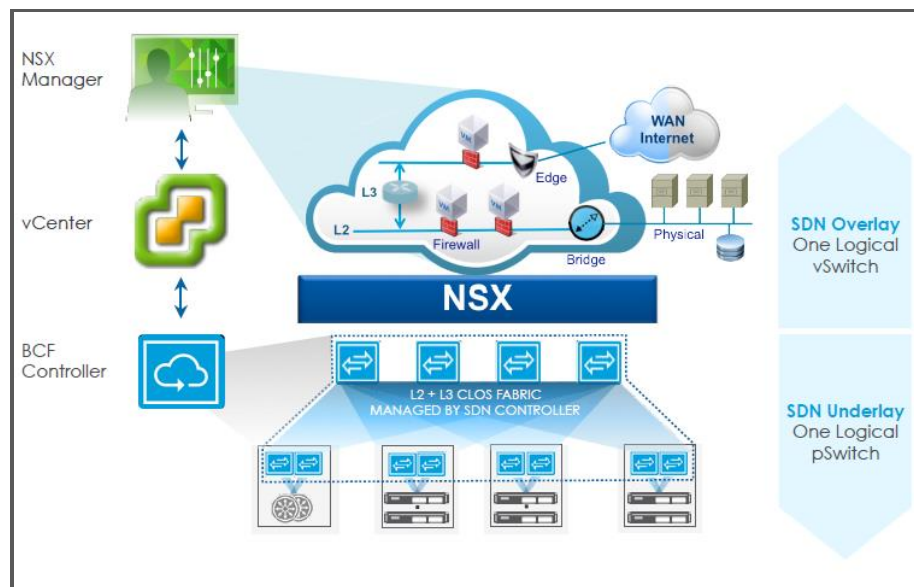


Figure 4. BCF Controller and vCenter Server Integration

Big Switch has implemented a vCenter extension in the BCF controller to exchange information about the VMware overlay network and create the correct relationships between the underlying physical fabric and the virtual overlay of the cloud. BCF uses LLDP and supporting controller automation to set up the associations between the underlay and overlay networks. This process is sometimes called Auto-LAG Formation. When comparing the BCF process to individually configuring the interfaces in a PMO leaf network of the same size the time savings are on the order of 18:1 in a fully configured POD. In the PMO an administrator needs to configure 40 individual interfaces per leaf switch toward the servers in its rack with link aggregation and protocol parameters. This process needs to be repeated across 32 leaves in the

16 switch POD. By contrast with its use of auto-discovery and auto-configuration processes, **BCF instantiates the underlay-to-overlay relationship in a fraction of the time, between 15-18 x as fast as in the PMO, freeing up cycles to focus on other goals.**

Five-Year Cumulative Total Cost of Operation

Combining capex and opex into a cumulative TCO for the VMware POD, **we find the five-year cumulative cost of the BCF solution to be approximately 40 percent of the cumulative TCO of PMO as shown in Figure 5 and Table 1.** With an average rate of acceleration in excess of 10:1 in key workflow tasks, the impact of the cost advantage for overall productivity and time to results is overwhelmingly clear.




Figure 5. Cumulative TCO of BCF and PMO in the VMware POD

	BCF-VMware	PMO-VMware
TOTAL	\$ 887,999	\$ 2,208,927

Table 1. Cumulative TCO of BCF and PMO in the VMware POD

BCF and PMO in an OpenStack-Based Private Cloud

With OpenStack BCF 3.0 brings a new dimension of comparison with PMO to bear in that BCF 3.0 includes Big Switch elements in both physical and virtual domains, constructing a unified P+V fabric. BCF Switch Light software is implemented in both the physical and virtual nodes and managed in unison by the BCF controller. This configuration is shown in the diagram in Figure 6.

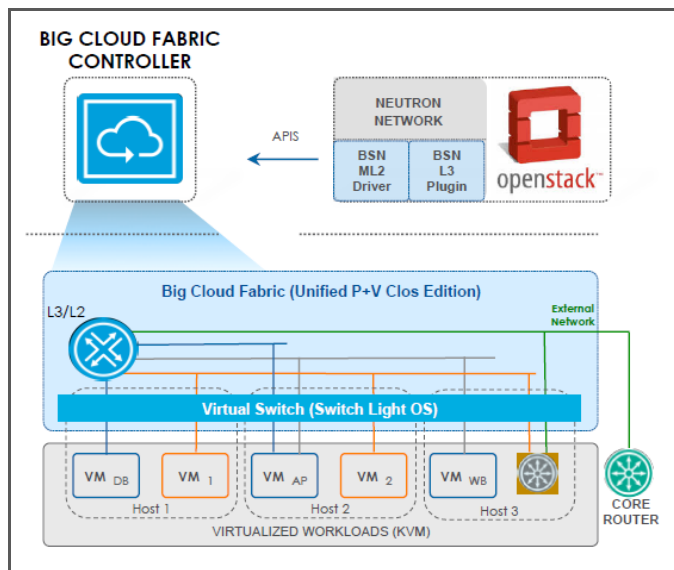


Figure 6. BCF Controller and OpenStack Server Integration

This implementation provides additional efficiency benefits to the OpenStack POD than were present in the VMware case, which relied on use of VMware’s virtual network overlays. Examples of these benefits are described in the opex cost comparison section.

Capex Differences in the OpenStack POD

As with the VMware case there are dramatic capital expense advantages of using BCF in the OpenStack POD. These continue to be based on the different hardware foundations of physical network switches of each alternative as well as significant differences in costs for support. But in the OpenStack case there is an additional benefit based in the single logical P+V fabric design of BCF versus the design for OpenStack using a PMO underlay. In the PMO case an additional overlay network is needed to support the operation of OpenStack virtual machines and to interface the OpenStack workloads with the PMO physical underlay. This is because OpenStack does not typically come with an overlay network solution of its own. Thus, an additional component is needed that adds its own capital and operating costs to the PMO case. Although each overlay virtual component has its own pricing and attributes, the fact that BCF is deployed and operated as a unified logical fabric versus PMO in which the components are separate adjacent domains produces additional differences in the efficiency of the choices.




Figure 7. Capex of BCF in the OpenStack POD

Capex for the BCF solution in the five-year window of the analysis is 35.8 percent of the PMO. Although the difference is list price based and discounting practices will vary between vendors that might comprise a PMO alternative in the general case, the difference is dramatic.




Figure 8. Capex of PMO in the OpenStack POD

OpenStack Opex Example: Path Trace Automation in the OpenStack Cloud. As in the VMware POD, here we highlight one area of opex efficiencies in which BCF produces significant efficiencies compared to the PMO. As in the VMware case there are others that contribute to the overall TCO advantage, but we highlight this one to bring it into clear focus in the comparison.

A significant requirement in cloud configurations is for users and operations teams to understand the operating status of the workloads users are running and the virtual machines (VM) on which they rely to do the work. Many times a query about the status of a workload or a VM involves checking to see whether the network path between VMs participating in an application design (say, in a three-tier application) is healthy. How one tests the status and health of a network path that includes both virtual and physical forwarding nodes between the VMs is a meaningful question in terms of how efficient a network solution for the cloud really is. If a CLI based, node-by-node query is required to be manually performed across five or six hops in a VM to VM path involving different network platforms at both P + V levels (versus a unified logical P + V fabric), it is easy to see where a PMO design would be cumbersome and inefficient.

On top of leveraging the efficiency of its own logically, centralized fabric design, Big Switch has embedded a function for automated access to path trace capabilities directly into the OpenStack dashboard to allow users of the cloud, such as application developers, to investigate VM to VM path availability rapidly. A view of the functionality is shown in Figure 9.

First, a view of how the visibility into the underlying fabric is made available for tests is shown. Tests are defined as those that work on simply defined descriptions of the application or group the user wants to evaluate. Indications of recent results and a set of choices for actions to perform now are presented.

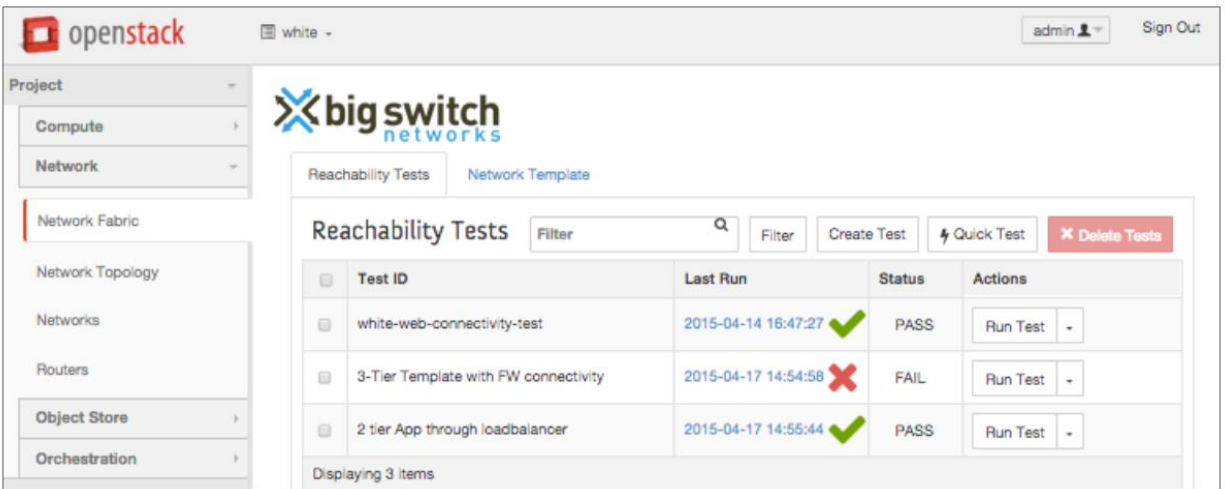


Figure 9. BCF Test Path Integration with OpenStack Horizon GUI

On execution of a test the results are presented in multiple ways, some of which are shown in Figure 10: the result (pass/fail), the path of elements traversed in the test (lower right), and other details of the instance involved are presented.

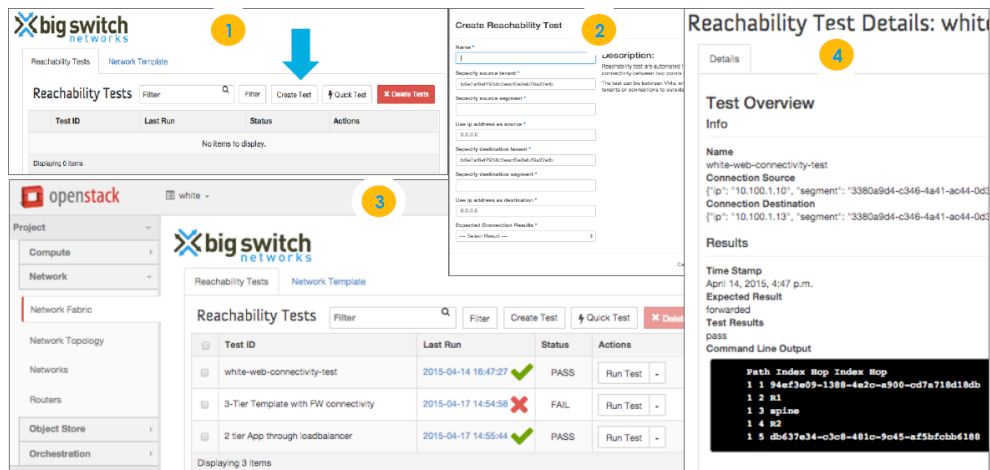


Figure 10. BCF Test Path Results in the OpenStack GUI

On average, automated path trace function in BCF is on the order of 12 times faster than the process required to obtain similar information in the PMO. Because the efficiencies are achievable by both operations and cloud end-user staff, the benefits accrue at a broader level than the PMO case in which the challenges of getting the answer rest fully in the operations team's hands. If we consider that these questions would arise on the order of five or six times per day in either case in a cloud configuration of this size (320 servers and well over a 1,000 VMs) the differences in efficiency of the affected staffs over a five-year cycle of operation are quite significant, as shown in Table 2. It is a true example of how the logically centralized and agile model of the BCF design is materially beneficial to its customer.

Diagnosing Problems + Evaluating Performance	VM-to-VM Path Trace Automation	\$ 26,041.67	\$ 312,500.00
	Self-service path tracing	\$ 26,041.67	\$ 312,500.00

Table 2. Opex Efficiencies from Test Path Automation with BCF

Cumulative Five-Year TCO. Combining the capex and opex measurements of the BCF and PMO deployments for the OpenStack cloud, we see the unification of P + V components in the single logical fabric producing even more amplified benefits than in the VMware case. **BCF is 29.6 percent of the TCO of the PMO design in this OpenStack POD (\$840,731 versus \$2,837,603).** Notably, much of the difference stems from the combination of software and opex efficiencies that in the end are the essence of the open SDN fabric vision.




Figure 11. Cumulative TCO of BCF and PMO in the OpenStack POD

Additional Benefits from Speed of Innovation and Accelerating Time to Results

In analyzing the two POD cases we identified the significant **TCO** advantages BCF provides in the private cloud. In addition to TCO BCF and its software create additional opportunities for customers to gain from innovations they might find to their advantage compared to the speed at which comparable innovations could be achieved in the PMO. In this manner, BCF becomes a key part of enabling differentiation and time to results—whether in revenue-generating opportunities from applications or increased efficiencies in groups beyond IT—and thus becomes an element in realizing business agility at a level not previously attainable.

A high-level view of this dynamic is conveyed in Figure 12, which illustrates revenue and profits being achieved at earlier milestones and higher levels when fundamental improvements in efficiency and pace of innovation are introduced into a firm's operating model. Techniques that lower development and operating costs and simplify innovation make direct contributions to improved results.




Figure 12. Accelerated Time to Results from Improved Innovation-Driven Fundamentals

In the case of private clouds we looked at two types of innovations firms might pursue that have a direct relationship to the network supporting their cloud and compare how the innovations would be achieved in the BCF and PMO alternatives. The results have a bearing on how efficiently the firm would be able to introduce the enhanced functions and how quickly it would start realizing the benefits of the innovations.

The first case we examined was how readily a firm could introduce a custom software enhancement to a cloud application being delivered that leverages fabric intelligence to its benefit – perhaps an enhancement to the criteria used for filtering users’ traffic for delivery to different tiers or modules in a crucial revenue-producing application (from e-commerce clusters, for example). This might involve making an enhancement to the security, forwarding, and chaining profiles that would apply to users accessing the modules. The organization’s own enhancement would be developed and integrated into the logic controlling the fabric. The faster and more easily it would be achievable, the more quickly the organization would be able to deliver its more refined customer experience.

To make this user-based enhancement feasible would involve developing software that leverages the REST APIs and higher level data models of the BCF controller versus developing substantially less integrated and more complex enhancements in a scripting language usable with the PMO. Our assessment is that **on average the development would be achievable on the order of four times more efficiently (combination of less staff required for development and less overall time required to implement based on efficiencies in coding techniques) for this class of enhancement versus the more cumbersome PMO automation approach.** If one were to envision making such enhancements several times over the course of five years with the cloud, then the benefits of more efficient innovation cycles would accumulate in a manner such as TCO does. If the benefits are improved user service, more targeted revenue offers, and increased efficiency in sales, then achieving those upsides earlier fits into the accelerated time to results profile previously outlined. The example in-house innovations we examined produce an example pattern of achievements as outlined in Table 3.

Service Innovation Advantages of BCF vs. PMO								
ID	Category Name	Scenario	Year 1	Year 2	Year 3	Year 4	Year 5	Total
2	Value of 'Own' New Functions Created In-House for Fabric	BCF-OpenStack	\$ 1,000,000	\$ 3,000,000	\$ 5,000,000	\$ 7,000,000	\$ 9,000,000	\$ 25,000,000
		PMO-OpenStack	\$ 250,000	\$ 1,250,000	\$ 2,250,000	\$ 3,250,000	\$ 4,250,000	\$ 11,250,000

Table 3. Example Advantages of BCF vs PMO for Customer-Driven In-House Innovations

The net effect of the comparison is that higher levels of benefit are achievable at earlier points in the deployment cycle when development processes are based on the BCF architecture versus the PMO. For an environment in which differentiated services produce meaningful improvements in customers' experiences or competitive position, this advantage is every bit as important as superior operating efficiencies via lower overall TCO when charting a course of operations.

In parallel with these innovation cycles, which are focused on a customer's ability to add value on their own, similar benefits can accrue from the agility each vendor has in making feature enhancements in their platform at customers' requests based on how flexible the architecture of the solution is and how efficiently new functional capabilities can be introduced. In the BCF case where the physical switching infrastructure is based on an open white box hardware design and where the switching, control, and enhanced applications software is based on using open APIs between processes, the difficulty of adding new functional enhancements (such as forwarding protocol extensions or integrations with unique application or orchestration functions) is substantially reduced compared to the less agile approach of modifying custom hardware or evolving less flexibly implemented software to accomplish the enhancement. BCF as a platform for evolution presents advantages similar to those it presents as a platform for innovation and a platform for operations.

The Evidence of a Turning Point is 'In'

When looking at bringing cloud agility to the networks clouds are relying on it is clear from our analysis that the benefits of a logically centralized, software-driven, unified P + V fabric provide a sustainable and dramatic advantage to customers compared to the cumbersome, more distributed and fragmented model of prior network implementations. 10:1 orders of magnitude in operations task efficiencies are simply too compelling to ignore. And when adding in the benefits of open fabric architectures in both hardware and software dimensions, with the overall benefits producing 60% and 70% improvements in total ownership costs, the appeal of the new direction is only amplified.

Based on these results we conclude that the evidence of a turning point is in, and that unified P+V, logically centralized, open SDN fabrics for data center networking are a design that has arrived, and are only likely to increase in their appeal in coming years.

ACG Research is an analyst and consulting company that focuses in the networking and telecom space. We offer comprehensive, high-quality, end-to-end business consulting and syndicated research services. Copyright © 2015 ACG Research.