Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING

WHITE PAPER

Making Networks SDN-Ready With Segment Routing

A Heavy Reading white paper produced for Cisco Systems Inc.

CISCO™

AUTHOR: STERLING PERRIN, PRINCIPAL ANALYST, HEAVY READING

OPERATOR REQUIREMENTS IN THE CLOUD ERA

Network Requirements

The era of the cloud has certainly arrived. According to Cisco's Global Cloud Index, which tracks data center IP traffic by cloud data center and traditional data center (i.e., not cloud), cloud accounted for two thirds of data center IP traffic in 2015, and is expected to rise to 80 percent of data center IP by 2019. The trend was initially driven by the Webscale Internet companies (Google, Facebook, Amazon, etc.), but today nearly everyone is adopting cloud – including colocation/carrier neutral providers, traditional telecom operators and enterprises.

The migration to cloud is leading to massive changes in how communications networks are built and operated. The key requirements for network operators in the cloud era include:

- **Capacity Scale:** We are seeing a large migration in metro networks now, from historical 10 Gbit/s rates to 100 Gbit/s, and now also 200 Gbit/s (using advanced modulation formats, 16 QAM). The primary driver for this is connecting data centers to other data centers, or data center interconnection (DCI).
- **Network and service agility:** The cloud model is based on sharing storage and computing resources across geographies with automation and on-demand. In order for these resources to be shared efficiently, however, the underlying communications network has to be both dynamic and flexible – a dramatic change from the static-pipes communications model of the past. Rapid re-configurability and automation need to be brought into the communications network itself.
- **Openness:** Operators have concluded that communications networks that respond and scale rapidly must become open in a way that has never before existed in telecommunications. Interoperability is required across domains, layers and vendors. To achieve this interoperability, open standards are needed; proprietary protocols and processes are simply too slow in the cloud era.

The Promise of SDN & Network Virtualization

Given the fundamental changes in network requirements, it is no surprise that the communications industry has identified two fundamentally new technology trends as the primary means to address them: software-defined networking (SDN) and network functions virtualization (NFV).

The following quote from a one-on-one interview conducted with a Tier 1 North American network operator in 2016 summarizes the values of SDN and virtualization succinctly: "We see both [SDN and NFV] as the keys to making the network automated and programmable. SDN is not the goal, and NFV is not the goal. The goal is a flexible, automated and programmable network to reduce opex and delivery services faster to market. That is the goal, and SDN and NFV are the tools to get there."

While early operator interest was primarily centered around cost reductions (including both capex and opex), goals have shifted over the past few years, and the latest Heavy Reading operator surveys point to rapid scaling and network and service agility as the primary drivers. Drawn from two separate operator surveys, **Figures 1** and **2** show expected operator benefits for SDN and NFV, respectively.

Figure 1: Most Important Business Driver for SDN

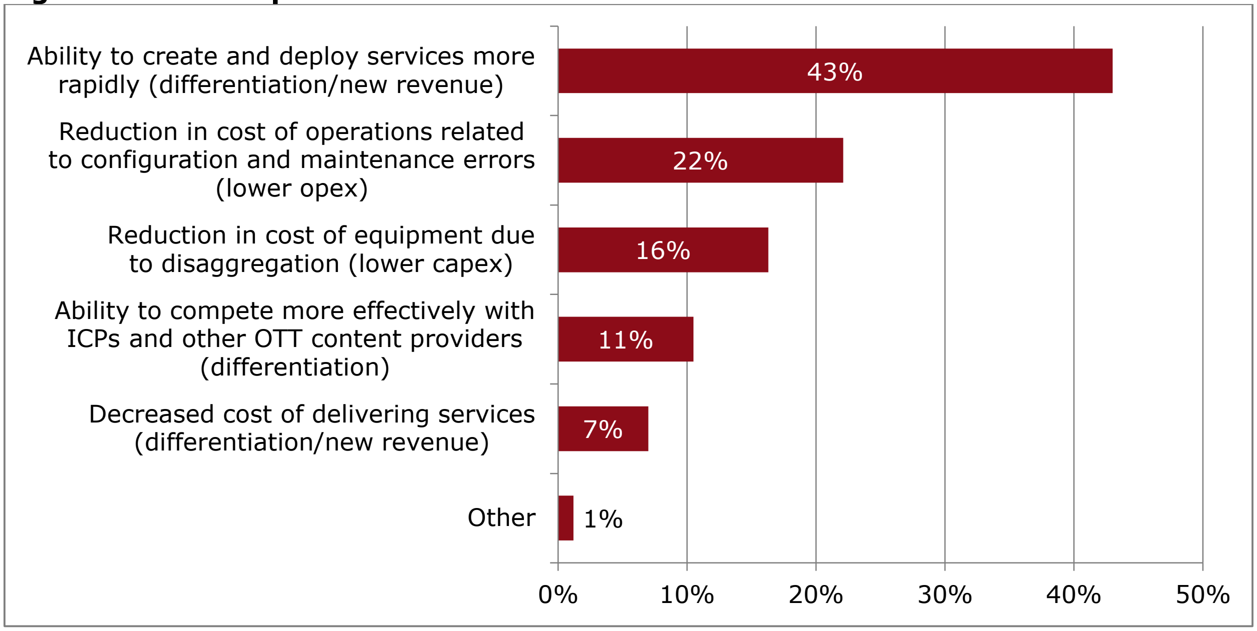

Source: "Carrier SDN: Service Provider Perspectives, Transition Strategies & Use Cases 2016: A Heavy Reading Multi-Client Study," June 2016; N=86

Figure 2: Most Important Expected Benefit of NFV

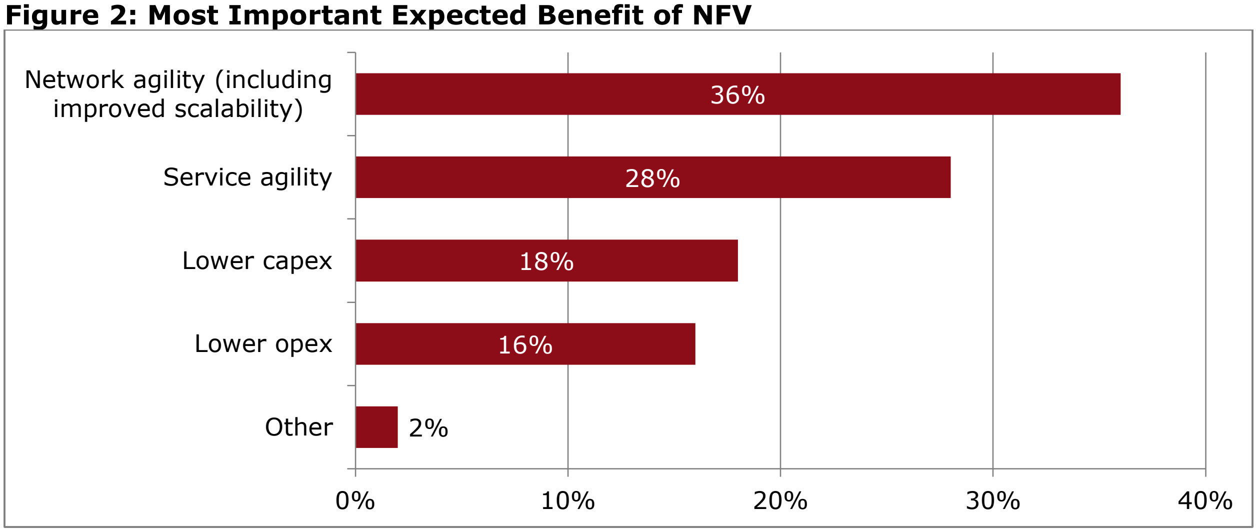

Source: Heavy Reading's May 2015 Network Transformation Survey, sponsored by Brocade; N=106

MAJOR NETWORK CHALLENGES IN THE SDN ERA

While the new era of SDN and virtualization is coming, it is still early days and many challenges remain. In a recent global survey of network operators published in June 2016, 74 percent of respondents reported that they are still in pre-commercialization phases of SDN deployment. Even for operators that have commercialized SDN, deployments tend to be limited in scope and footprint.

Network complexity is one of the reasons why deployments are limited even when SDN and virtualization are deployed commercially. The carrier WAN, for example, is far more complex than the data center environment and, for this reason, we have seen SDN expand in intra-DC deployments but stall when it hits the WAN. There is no question that intra-DC traffic volumes are massive, but traffic volume alone does not determine overall network complexity. Carrier networks have hundreds of thousands of nodes spanning wide geographies (national and global), many services and protocols, unpredictable traffic patterns, and millions of users.

The maturity of SDN in the WAN reflects this greater network complexity. Vendor SDN products targeting DCs have been commercially available for years and deployments are numerous, but carrier WAN SDN products are much newer to market, and commercial deployments are far fewer to date.

Differences in data center and carrier WAN networks – including protocols used and SDN availability/maturity – creates another challenge. Many operators see SDN as an opportunity to bridge together their data centers with the WAN, so that the needs of the data center-based applications can be appropriately and rapidly met by the network that connects the data centers and the users.

However, to date, data centers and networks have been on different trajectories and timetables, making it difficult for this efficient bridging of domains to take place. Yet, without tight coupling of the applications and the network, the benefits of virtualization and SDN are greatly reduced.

Software control is another challenge, and one that reflects the early phase of evolution. Centralized SDN control is a fundamental tenet of SDN, but how to best achieve that centralized control remains a topic of debate. Originally, all SDN was based on OpenFlow, so the control issue was simple: It was OpenFlow. However, enterprises, service providers, and network operators quickly realized that OpenFlow had many limitations when it came to the WAN (including scalability and resiliency) and began exploring alternative protocols.

Operators have tested centralized control in labs and isolated environments, but the move to large-scale commercial networks has been slow. Today, the best means of software control varies by provider and by application, and remains an open question.

The final challenge flows from the fact that operators did not select OpenFlow as the universal SDN protocol for all applications. There has, in fact, been a vast proliferation of protocol and standards propositions for SDN over the past four years, leading to a new challenge that Telefónica has called the "ocean of protocols" surrounding SDN.

Two problems have arisen from the ocean of protocols:

- **Operator confusion:** With too many choices, operators don't know the differences between the options or which technology (or technologies) is the right one. The end result has been for operators to extend their timelines for rollouts as they evaluate all these options.
- **Lack of interoperability and compatibility:** Protocols and architectures may perform well in isolation, but real-world networks consist of many protocols and systems. If the new technologies are incompatible with the installed base, many operators see little value in moving forward. Thus, ensuring compatibility (particularly with installed networks and technologies) is of critical importance to network operators.

SEGMENT ROUTING: AN INTRODUCTION

Against this backdrop of traffic demands and the coming SDN and virtualization era, **segment routing** has emerged as highly promising and practical solution.

Segment routing is becoming a popular topic today, but, it fact, it is not a new technology concept in IP routing. Segment routing is actually a variation of source routing, a routing technique in which the sending router specifies the route that the packet will take through the network, rather than the path being chosen based on the packet's destination only. Source routing as a concept has existed in research and academia for some 20 years, but implementation is relatively new.

In segment routing, a node steers a packet through an ordered list of instructions called "segments." A segment can represent any instruction, whether based on topology or service. As with other source routing techniques, the full instructions for the path through the network are embedded in the packet header, and this is applied at the source node. In this case, these are MPLS headers on IPv4 packets today (and directly on IPv6 packets in the future, as we will discuss later in the paper).

With segment routing, there are two different types of segments. Global segments route the traffic over the shortest path to the destination, as computed by the IGP (IS-IS/OSPF). Local segments are applied on a per-hop basis and used to divert traffic from the shortest path whenever desired (i.e., reasons of latency, redundancy, SLA requirements, etc.) To be clear, applying a new label to every hop is an option in segment routing but is not mandatory, due to the existence of global segments.

Instructions are applied to the packets as an ordered stack of labels. Each router processes the segment at the top of the stack, then removes the top segment and sends the packet according to the instruction. At the next hop, the new router processes the new top segment, removes it from the stack, and sends the packet on its way. This process is followed until all segments are removed, and the packets reach their destination.

Figure 3 provides a simplified diagram of segment routing through an MPLS network and includes the combination of both global segments and local segments. In the figure, segment information is embedded in the segment router header at the ingress router (R1), defining an explicit path to the destination router (R6) via intermediate routers R2 and R3. In order to reach R3 via R2, local segment labels must be applied (because it is not the shortest path). However, after reaching R3, no explicit path is defined, and the global segment is applied. In this example, two equal routes exist to reach the destination (R6). Therefore, equal-cost multi-path (ECMP) routing load-balances packets across the two equal routes to reach the destination.

Significantly, segment routing is not a proprietary concept and is, in fact, well on its way to becoming standardized through the IETF, through "draft-ietf-spring-segment-routing-10." The first IETF segment routing draft was published in 2013, and it is currently in the "last call" phase to become standardized. Draft contributors include authors from Cisco, Orange, Google Ericsson, Deutsche Telekom, and others.

Figure 3: Segment Routing Through an MPLS Network

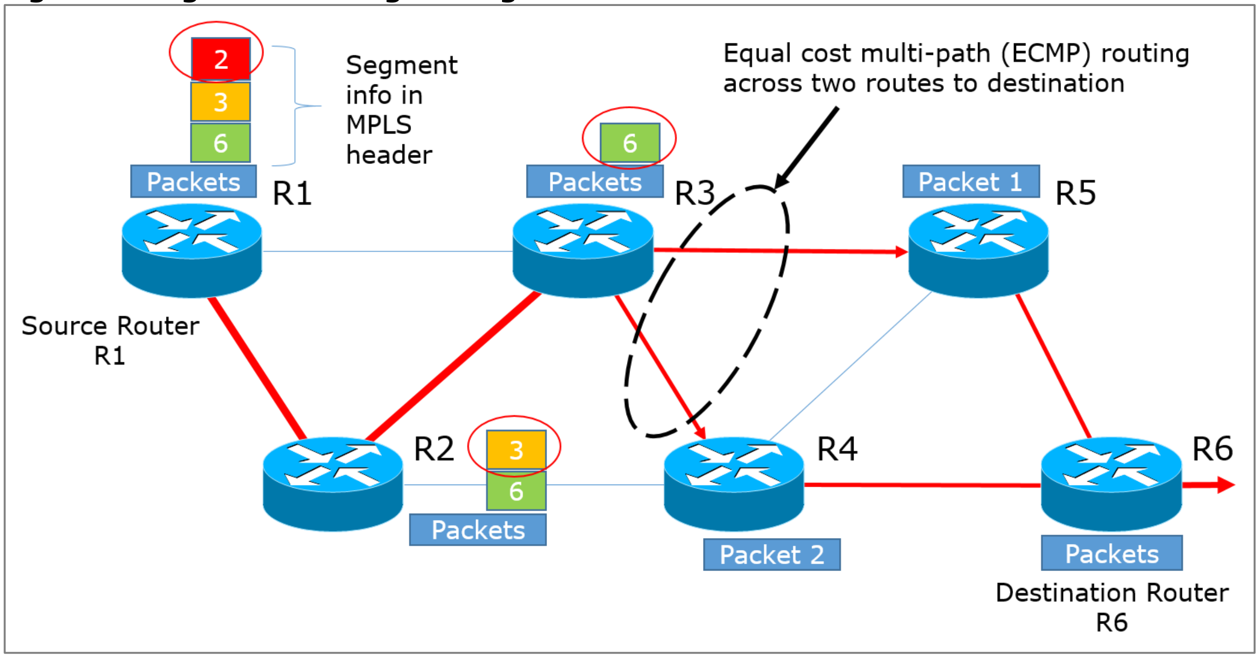

Source: Heavy Reading, 2016

ADDRESSING NETWORK CHALLENGES WITH SEGMENT ROUTING

Removing Complexity, Increasing Scale

Operators use traffic engineering (TE) to create efficient and reliable IP network operations, and to optimize network resource utilization and network performance. Key general benefits of traffic engineering include the following:

- **Minimize network congestion to avoid traffic blockage and boost network performance.**
- **Increase overall network efficiency.** While TE doesn't create new capacity, it allows existing capacity to achieve higher levels of utilization, thus delaying the need for additional capacity and reducing network capex.
- **Opex reduction through automating traffic decisions**.
- **Defining recovery paths in case of failure to assure reliability and network uptime**. MPLS Fast Reroute is an example of this.
- **Defining class of service (CoS):** TE can be used to create different paths and priorities for different types of traffic, based on the priorities assigned to those traffic flows.

With traditional MPLS traffic engineering, using RSVP-TE, all of the routers along the engineered route must maintain state – meaning that these routers must be updated with information about the end-to-end path and nodes. Maintaining router state throughout a network, however, adds a great deal of complexity to the network's operation, and this complexity has greatly hindered network scalability using traffic engineering.

Segment routing, however, eliminates the scalability/complexity limitation by requiring that only the ingress router (i.e., R1 as illustrated in **Figure 3**) hold state data. All of the required state information for the end-to-end connection is contained in the segment list, so neither the transit routers nor the egress router needs to hold state information. In the simple four-router path example from **Figure 3**, only R1 holds the state information – a 75 percent reduction in state information/complexity compared with engineering using RSVP-TE. This is the single biggest reason why segment routing has improved scalability.

Works With SDN-Based Control

Segment routing does not require SDN control in the network, but the routing technology was created with centralized SDN control in mind, so that segment routing and SDN are complementary technologies. To be clear, operators can gain immediate benefits by using segment routing on distributed router networks, with no centralized SDN control. One distributed segment routing use case is operations and protocol simplification (compared to traditional MPLS). A second use case is 50 millisecond protection at Layer 3.

However, combining segment routing with SDN control expands the set of available use cases and allows operators to gain the maximum set of benefits from the source routing technology. For the majority of operators, SDN control is the end game – even if they are only using distributed control for their initial deployments. (The Bell Canada profile below provides one example of this phased approach.)

In particular, a centralized controller can actively collect and monitor topology and traffic-engineering state changes with a global view of the network. Path computation element protocol (PCEP) can then be used to present near-real-time views of topology deployment. The PCE can be contained within a vendor's SDN controller, through which network commands can be given and acted upon. **Figure 4** illustrates the centralized collection and reporting abilities of a PCE coupled with segment routing-based traffic engineering on an end-to-end network.

Figure 4: Centralized Control With Path Computation Element

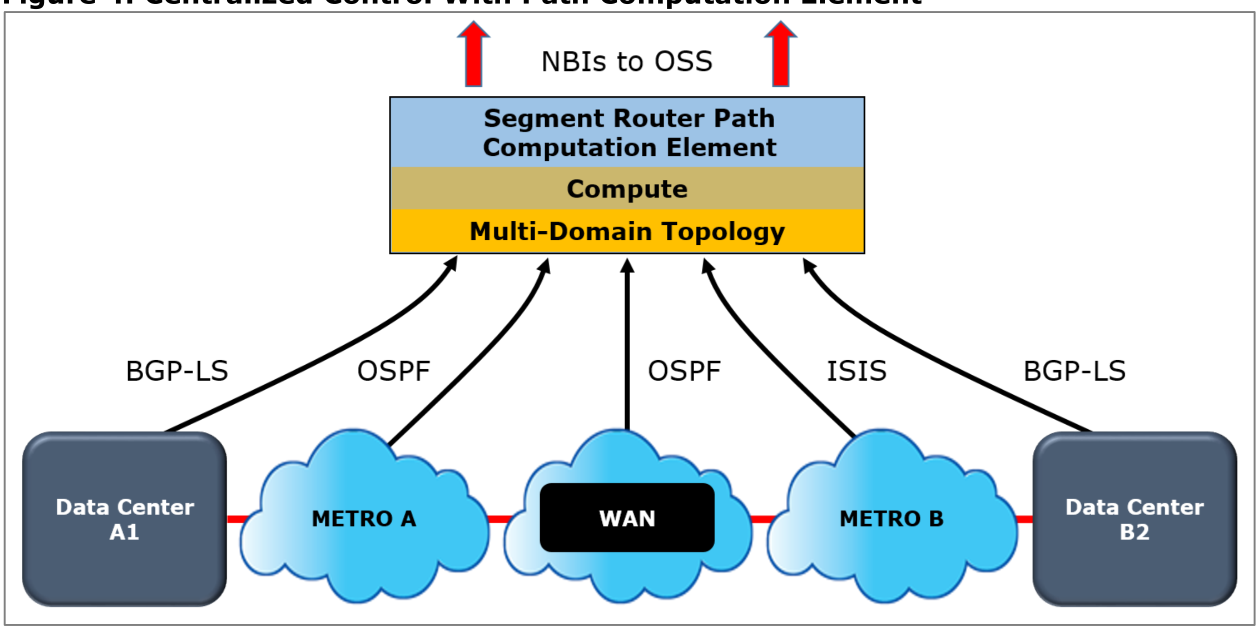

Source: Cisco and Heavy Reading, 2016

One benefit of using segment routing with SDN is that there are significant improvements in convergence times, due to the limited amount of state information that must be distributed by the SDN controllers – since, with segment routing, all required state information is created in the header at the ingress router.

Another benefit of using segment routing is interoperability across vendors and domains, and with existing networks. Segment routing developers deliberately chose to build with known protocols and not define a new control plane protocol. They also chose to define segment routing on MPLS networks, which are widely deployed by operators around the world (adding IPv6 support for the future).

The use of MPLS combined with existing protocols such as BGP, PCEP, NETCONF/YANG, etc., means that segment routing can easily be inserted into – and interoperate with – existing IP networks. Furthermore, segment routing is being standardized within the IETF, ensuring multi-vendor support.

As a final significant point, the combination of segment routing with centralized PCE allows traffic engineered paths to be defined across multiple domains – defining paths from metro networks to core networks, but also connecting data centers to the WAN. While some other source routing techniques are coupled tightly with OpenFlow-based SDN, these techniques are ultimately limited to operate only where OpenFlow is present in the network. With segment routing as defined by the IETF, this limitation does not exist.

BELL CANADA: OPERATOR USE CASE

Bell Canada is a Tier 1 network operator with a long history and a large installed base of legacy networks across both wireline and wireless. Its current network transformation project is called Network 3.0, and its goals are to deliver the best experience to its customers via new software-driven and cloud technologies.

Bell Canada lists its next-generation network requirements as follows:

- Needs to be an industry standard, ratified by global standards organizations
- Reusable in the core/WAN, possibly as the glue to bring all the networks together
- Software-programmable
- Leverage new CO/DC greenfield opportunity to try something new
- Provides solutions for both transition and end state
- Interoperability with both the brownfield and greenfield
- Implicit ECMP handling

For Bell Canada, segment routing meets these NGN requirements, and it is being adopted as part of its phased Network 3.0 transformation project. The operator's end goal is a combination of segment routing and centralized SDN control to enable on-demand traffic engineering end-to-end across its networks, from the data center to the WAN, and from the access network to the core.

On-demand traffic engineering provides two key benefits:

- **Increasing network utilization:** Simply adding capacity to the network is no longer sufficient to meet traffic demands, because the revenue derived per bit cannot cover the costs per bit. Greater efficiency is also required, so that operators get the most from every link in the network. Bell Canada calls this the ability to mine "sleeping capex."
- **Drastically simplifying network operations**: Traffic engineering in MPLS networks is manual, complex and time-consuming. Segment routing simplifies the TE function, while also eliminating months of upfront planning. The results are reduced opex and greater network agility, which is aligned with what cloud and virtualization require.

Figure 5 illustrates the end of application-responsive network end-to-end across the Bell Canada network.

Figure 5: Illustration of End-to-End Segment Routed Network

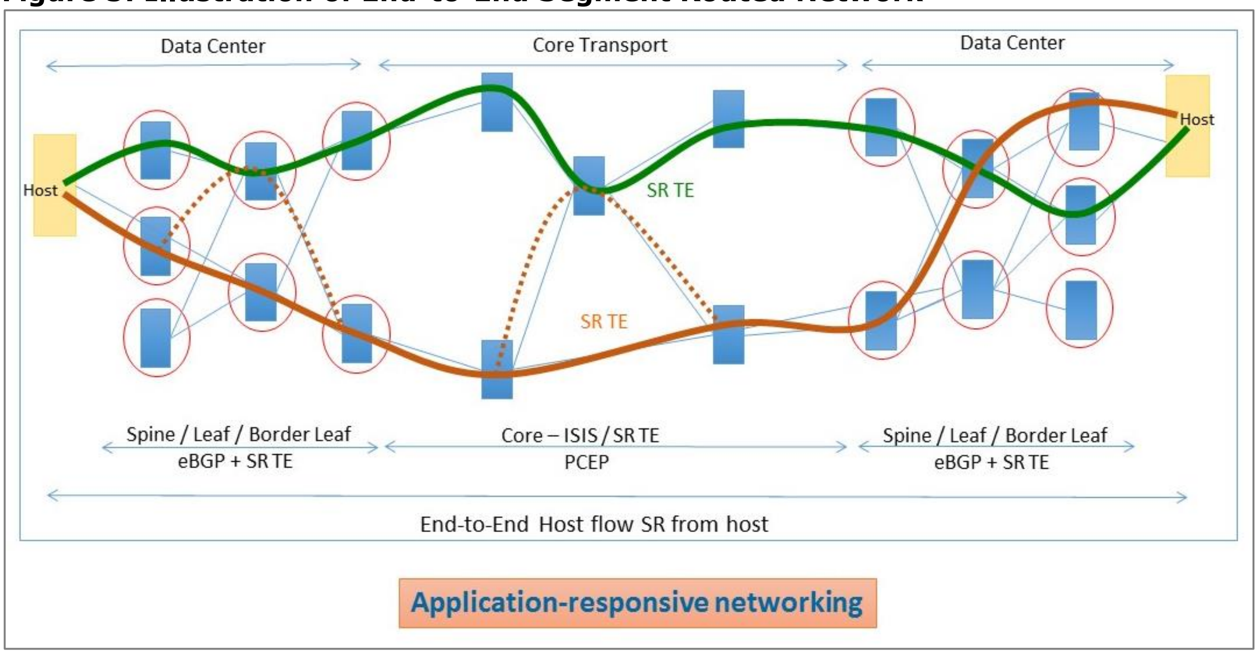

Source: Bell Canada, 2016

Consistent with its NGN requirements to maintain both greenfield and brownfield networks, Bell Canada is adopting a phased approach to segment routing. Step 1 is to retain the distributed MPLS network but to start adding segment routing, so that the architecture, engineering and operations teams can start getting used to the new technology. Step 2 is to begin introducing centralized control, via PCEP, on certain "islands" in the network to begin testing features. Step 3 will be to roll out segment routing across the Bell Canada infrastructure.

The final steps are longer-term. With segment routing established in the network, the operator will be able to simplify its protocol stack by removing LDP and RSVP-TE. (Segment routing will have taken over their job). Next will come the introduction of IPv6 into the network, with segment routing running on the new IPv6 and legacy MPLS networks. The final, longer-term phase will be the elimination of MPLS entirely.

CONCLUSIONS & NEXT STEPS

Although not new in concept, segment routing is becoming a popular topic today, as a key means of scaling capacity and creating network agility in the new era of virtualization and cloud services delivery. Key to its value proposition is its ability to automate and vastly simplify traffic engineering compared to current MPLS-based networks. With traffic engineering simplified and on-demand, TE goals can be achieved including:

- Increased network efficiency (reducing capex)
- Boosted network performance
- Increased opex efficiency through automation (reduced opex and faster operations)
- Assured reliability and network uptime
- End-to-end class-of-service (CoS) assurance

Segment routing and SDN are highly complementary. Significantly, segment routing was designed with SDN in mind but does not require SDN to be effective. With a phased approach, operators can simplify their networks using segment routing with distributed control on Day 1 and can evolve to centralized SDN control over time. Achieving the full set of benefits with segment routing assumes centralized SDN control and is the end goal for the majority of operators.

With IETF standardization in place and SDN technologies maturing (i.e., PCEP), Heavy Reading expects segment routing adoption to accelerate globally among operators of all sizes. The phased evolution – as outlined in the Bell Canada profile – provides a solid blueprint for Tier 1 and Tier 2 operators to follow as they design for tomorrow while bridging to the networks of today. Ultimately, the IP world will migrate from IPv4 to IPv6, and the segment routing standard has planned for this too. Operators that begin adopting segment routing today will be ready.