

Chapter 7

IMPLEMENTING NOVEL DEFENSE FUNCTIONALITY IN MPLS NETWORKS USING HYPERSPEED SIGNALING

Daniel Guernsey, Mason Rice and Sujeet Shenoj

Abstract Imagine if a network administrator had powers like the superhero Flash – perceived invisibility, omnipresence and superior surveillance and reconnaissance abilities – that would enable the administrator to send early warnings of threats and trigger mitigation efforts before malicious traffic reaches its target.

This paper describes the hyperspeed signaling paradigm, which can endow a network administrator with Flash-like superpowers. Hyperspeed signaling uses optimal (hyperspeed) paths to transmit high priority traffic while other traffic is sent along suboptimal (slower) paths. Slowing the traffic ever so slightly enables the faster command and control messages to implement sophisticated network defense mechanisms. The defense techniques enabled by hyperspeed signaling include distributed filtering, teleporting packets, quarantining network devices, tagging and tracking suspicious packets, projecting holographic network topologies and transfiguring networks. The paper also discusses the principal challenges involved in implementing hyperspeed signaling in MPLS networks.

Keywords: Hyperspeed signaling, network defense, MPLS networks

1. Introduction

The midnight ride of Paul Revere on April 18, 1775 alerted the Revolutionary Forces about the movements of the British military before the Battles of Lexington and Concord. The ability to deploy Paul-Revere-like sentinel messages within a computer network could help improve defensive postures. These sentinel messages could outrun malicious traffic, provide early warnings of threats and trigger mitigation efforts. Electrons cannot be made to move faster than the laws of physics permit, but “suspicious” traffic can be slowed down ever so

slightly to enable sentinel messages to accomplish their task. To use an optical analogy, it is not possible to travel faster than light, but “hyperspeed signaling paths” can be created by slowing light along all other paths by increasing the refractive index of the transmission media.

The concept of offering different priorities – or speeds – for communications is not new. The U.S. Postal Service has numerous classes of mail services ranging from ground delivery to Express Mail that guarantees overnight delivery. The U.S. military’s Defense Switched Network (DSN) [9] designed during the Cold War had four levels of urgency for telephone calls, where a call at a higher level could preempt a call at a lower level; the highest level was FLASH, which also incorporated a special FLASH OVERRIDE feature for the President, Secretary of Defense and other key leaders during defensive emergencies. Modern MPLS networks used by major service providers offer a variety of high-speed and low-speed paths for customer traffic based on service level agreements.

This paper proposes the use of optimal (hyperspeed) paths for command and control (and other high priority) traffic and suboptimal (slower) paths for all other traffic in order to implement sophisticated network defense techniques. The basic idea is to offer a guaranteed reaction time window so that packets sent along hyperspeed paths can arrive sufficiently in advance of malicious traffic in order to alert network devices and initiate defensive actions. Separate channels have been used for command and control signals. Signaling System 7 (SS7) telephone networks provide a back-end private network for call control and traffic management, which physically separates the control and data planes [13]. MPLS networks logically separate the internal IP control network from external IP networks that connect with the data plane [7].

This paper describes the hyperspeed signaling paradigm, including its core capabilities and implementation requirements for MPLS networks. Novel defense techniques enabled by hyperspeed signaling, ranging from distributed filtering and teleportation to quarantining and network holography, are highlighted. The paper also discusses the principal challenges involved in implementing hyperspeed signaling, which include network deployment, traffic burden and net neutrality.

2. Hyperspeed Signaling

Hyperspeed signaling uses optimal (hyperspeed) paths to transmit high priority traffic; other traffic is sent along suboptimal (slower) paths. The difference in the time taken by a packet to traverse a hyperspeed path compared with a slower path creates a reaction time window that can be leveraged for network defense and other applications. Indeed, a hyperspeed signaling path between two network nodes essentially induces a “quantum entanglement” of the two nodes, allowing them to interact with each other seemingly instantaneously.

In general, there would be one or more hyperspeed (optimal) paths and multiple slower (suboptimal) paths between two nodes. Thus, different reaction time windows would be available for a hyperspeed path compared with (different) slower paths. These different windows would provide varying amounts

of time to implement defensive actions. Depending on its nature and priority, traffic could be sent along different suboptimal paths. For example, traffic deemed to be “suspicious” could be sent along the slowest paths.

Note that hyperspeed paths need not be reserved only for command and control traffic. Certain time-critical traffic, such as interactive voice and video communications, could also be sent along faster, and possibly, hyperspeed paths. Of course, using faster paths for all traffic would reduce the reaction time windows, and, consequently, decrease the time available to implement defensive actions. Clearly, a service provider or network administrator would prefer not to reduce traffic speed drastically. Consequently, a suboptimal path should incorporate the smallest delay necessary to obtain the desired reaction time window.

3. Core Capabilities

Hyperspeed signaling provides the network administrator with “powers” like the superhero Flash. The reaction time window corresponds to the speed advantage that Flash has over a slower villain. The ability to send signals between two network nodes faster than other traffic provides superpowers such as perceived invisibility, omnipresence and superior intelligence, surveillance and reconnaissance abilities.

This section describes the core capabilities offered by hyperspeed signaling. These capabilities are described in terms of the “See-Think-Do” metaphor [15].

3.1 Omnipresence

Omnipresence in the hyperspeed signaling paradigm does not imply that the network administrator is everywhere in the network at every moment in time. Instead, omnipresence is defined with respect to a target packet – the network administrator can send a hyperspeed signal to any node in the network before the target packet arrives at the node.

Omnipresence with respect to multiple packets has two versions, one stronger and the other weaker. The stronger version corresponds to a situation where there is one Flash, and this Flash can arrive before all the packets under consideration arrive at their destinations. The weaker version corresponds to a situation where there are multiple Flashes, one for each packet under consideration. Note that the stronger version of omnipresence endows the network administrator with the ability to track multiple packets and to correlate information about all the packets regardless of their locations in the network.

3.2 Intelligence, Surveillance, Reconnaissance

Intelligence, surveillance and reconnaissance (ISR) are essential elements of U.S. defensive operations [4]. ISR capabilities are implemented in a wide variety of systems for acquiring and processing information needed by national security decision makers and military commanders.

Intelligence is strategic in nature; it involves the integration of time-sensitive information from all sources into concise, accurate and objective reports related to the threat situation. Reconnaissance, which is tactical in nature, refers to an effort or a mission to acquire information about a target, possibly a one-time endeavor. Surveillance lies between intelligence and reconnaissance. It refers to the systematic observation of a targeted area or group, usually over an extended time.

Obviously, hyperspeed signaling would significantly advance ISR capabilities in cyberspace. The scope and speed of ISR activities would depend on the degree of connectedness of network nodes via hyperspeed paths and the reaction time windows offered by the paths.

3.3 Defensive Actions

Hyperspeed signaling can help implement several sophisticated network defense techniques. The techniques resemble the “tricks” used in stage magic. In particular, the advance warning feature provided by hyperspeed signaling enables a network to seemingly employ “precognition” and react to an attack before it reaches the target. As described in Section 6, hyperspeed signaling enables distributed filtering, teleporting packets, quarantining network devices, tagging and tracking suspicious packets, projecting holographic network topologies, and transfiguring networks.

Distributed filtering permits detection mechanisms to be “outsourced” to various locations and/or organizations. Teleportation enables packets to be transported by “secret passageways” across a network without being detected. Quarantining enables a network device, segment or path to “vanish” before it can be affected by an attack. Tagging facilitates the tracking of suspicious traffic and routing other traffic accordingly. Network holography employs “smoke and mirrors” to conceal a real network and project an illusory topology. Transfiguration enables network topologies to be dynamically manipulated (i.e., “shape shifted”) to adapt to the environment and context.

4. Multiprotocol Label Switching Networks

Circuit switching and packet switching are the two main paradigms for transporting traffic across large networks [10]. ATM and Frame Relay (OSI Layer 2) are examples of circuit-switched (i.e., connection-oriented) technologies that provide low latency and high quality of service (QoS). IP (OSI Layer 3) is a packet-switched (i.e., connectionless) protocol that unifies heterogeneous network technologies to support numerous Internet applications.

An important goal of service providers is to design networks with the flexibility of IP and the speed of circuit switching without sacrificing efficiency [8]. In traditional implementations, an overlay model is used, for example, to create an ATM virtual circuit between each pair of IP routers. Operating independently, the two technologies create a relatively inefficient solution. Since IP routers are unaware of the ATM infrastructure and ATM switches are unaware of IP

routing, an ATM network must present a virtual topology such as a complete mesh (which is expensive) or a hub with spokes (which lacks redundancy) that connects each IP router. IP may then be used to route traffic based on the virtual, rather than physical, topology.

On the other hand, the routing control paradigm used in IP networks is closely tied to the forwarding mechanism. Since a classless inter-domain routing (CIDR) IP address consists of a network prefix followed by a host identifier, IP networks have a hierarchical model. IP nodes forward packets according to the most specific (“longest match”) route entry identified by the destination address. Consequently, IP networks are only compatible with control paradigms that create hierarchical routes.

The need to enhance QoS and integrate IP with connection-oriented technologies like ATM has prompted the development of a more general forwarding scheme for MPLS – one that does not limit the control paradigm [5, 7]. This forwarding mechanism, called “label switching,” is similar to the technique used by circuit-switched technologies. Thus, MPLS enables connection-oriented nodes to peer directly with connectionless technologies by transforming ATM switches into IP routers. ATM switches participate directly in IP routing protocols (e.g., RIP and OSPF) to construct label switched paths (LSPs). LSPs are implemented in ATM switches as virtual circuits, enabling existing ATM technology to support the MPLS forwarding mechanism. Conversely, MPLS enables connectionless technologies, e.g., Ethernet, to behave in a connection-oriented manner by augmenting IP addresses and routing protocols with relatively short, fixed-length labels.

MPLS employs a single adaptable forwarding algorithm that supports multiple control components. MPLS labels are conceptually similar to the bar codes on U.S. mail that encode ZIP+4 information; these bar codes are used by the U.S. Postal Service to automatically sort, prioritize, route and track nearly 750 million pieces of mail a day. Within the MPLS core, label switching relies only on the packet label to select an LSP. Thus, any algorithm that can construct LSPs and specify labels can be used to control an MPLS network core. Some additional components are required at the edge where the MPLS core connects to other types of networks (e.g., an inter-office VPN running traditional IP). The MPLS edge routers interpret external routing information, place labels on ingress packets and remove labels from egress packets.

The following sections describe label switching and label distribution, which underlie packet transport in MPLS networks.

4.1 Label Switching

MPLS packet forwarding resembles the mechanism used in circuit-switched technologies; in fact, it is compatible with ATM and Frame Relay networks [5, 7]. Each MPLS label is a 32-bit fixed-length tag that is inserted in the Layer 2 header (e.g., for ATM VCI and Frame Relay DLCI) or in a separate “shim” between Layers 2 and 3 [12]. A label works much like an IP address in that it dictates the path used by a router to forward the packet. Unlike an IP




Figure 1. MPLS packet forwarding.

address, however, an MPLS label only has local significance. When a router receives a labeled packet, the label informs the router (and only that router) about the operations to be performed on the packet. Typically, a router pops the label on an incoming packet and pushes a new label for the router at the next hop in the MPLS network; the network address in Layer 3 is unchanged.

MPLS networks carry traffic between other connected networks. As such, most user traffic travels from ingress to egress (i.e., the traffic is neither destined for nor originating from internal MPLS hosts). At the ingress, a label is placed in the packet between the OSI Layer 2 and 3 headers [12]. The label informs the next hop about the path, destination and relative “importance” of the packet. At each hop, the label is examined to determine the next hop and outgoing label for the packet. The packet is then relabeled and forwarded. This process continues until the packet reaches the egress where the label is removed. If the MPLS network is composed mainly of ATM switches, the ATM hardware can naturally implement the MPLS forwarding algorithm using the ATM header with little or no hardware modification.

Figure 1 shows a typical MPLS architecture that interconnects two customer VPN sites. Routers A through F in the MPLS network are called label switching routers (LSRs). Customer edge routers, CE1 and CE2, reside at the edge of the customer network and provide MPLS core connectivity.

Consider the label switched path (LSP) from VPN Site 1 to VPN Site 2 (Routers A, B, C and F). Router A is designated as the ingress and Router F is designated as the egress for the path. The ingress and egress nodes are called label edge routers (LERs) [12]. When an IP packet reaches the ingress of the MPLS network, LER A consults its forwarding information base (FIB) and assigns the packet to a forwarding equivalence class (FEC). The FEC maps to a designated label that specifies QoS and class of service (CoS) requirements based on Layer 3 parameters in the packet (e.g., source IP, destination IP and application ports). In this example, LER A pushes Label L1 onto the packet and forwards it to LSR B. LSR B reads the label and consults its incoming label map (ILM) to identify the FEC of the packet. It then pops the previous label, pushes a new label (L2) and forwards the packet to its next hop LSR C. LSR C behaves similarly, forwarding the packet to LER F. LER F then pops

L3, examines the destination IP address and forwards the packet to VPN Site 2, where traditional IP forwarding resumes.

4.2 Label Distribution

A forwarding algorithm alone is not enough to implement an MPLS network. The individual nodes need to know the network topology in order to make informed forwarding decisions. The MPLS Working Group [1] defines a forwarding mechanism and control components to emulate IP routes using MPLS labels and paths. In IP, routing protocols such as RIP and OSPF populate IP forwarding tables [10]. Similarly, MPLS requires label distribution protocols to build end-to-end LSPs by populating the FIB and ILM of each node. Because MPLS is not tied to a particular paradigm, any routing protocol capable of carrying MPLS labels can be used to build MPLS LSPs. Such protocols include:

- **Label Distribution Protocol (LDP):** This protocol is designed to build aggregate LSPs based on IP routing information gathered by a traditional IP routing protocol such as RIP [1].
- **Resource Reservation Protocol – Traffic Engineering (RSVP-TE):** This protocol incorporates extensions to RSVP in order to construct LSP tunnels along requested paths with varying QoS. RSVP-TE is commonly used for traffic engineering (TE) in MPLS networks [2].
- **Multiprotocol Extension to Border Gateway Protocol 4 (MP-BGP):** This protocol extends BGP, and generalizes distributed gateway addresses and carries labels. It is commonly used to build VPNs [3, 11].

The three protocols listed above are commonly employed in IP-based networks. This demonstrates that MPLS seamlessly supports the IP routing paradigm and enables IP to be efficiently deployed in ATM and Frame Relay networks without the need for convoluted virtual topologies.

5. MPLS Implementation Requirements

Two requirements must be met to implement hyperspeed signaling. First, the network must be able to recognize and distinguish hyperspeed signals from non-hyperspeed signals. Second, the network must be able to provide appreciable differences in delivery delays, so that the reaction time windows are satisfied by hyperspeed signals. The network environment and the delay techniques that are applied govern the degree of flexibility with respect to the maximum possible reaction time window.

MPLS is an ideal technology for implementing hyperspeed signaling because it has built-in identification and service differentiation technologies. Labels in MPLS act like circuit identifiers in ATM to designate the paths taken by packets in the network core.




Figure 2. Egress filtering.

Hyperspeed signaling in MPLS would use the packet label to distinguish hyperspeed packets from non-hyperspeed packets. MPLS-capable routers are typically equipped with many QoS and traffic shaping features. LSRs can be configured to give hyperspeed packets the highest priority based on the packet label. Likewise, LSRs can be configured to delay non-hyperspeed packets in forwarding queues. Because the label dictates the QoS and the path, non-hyperspeed packets could be forced to take circuitous routes by constructing the corresponding LSPs using non-hyperspeed labels. The labels corresponding to optimal routes are reserved for hyperspeed packets.

6. Novel Defense Techniques

Hyperspeed signaling can help implement several sophisticated network defense techniques. These include distributed filtering, teleporting packets, quarantining network devices, tagging and tracking suspicious packets, projecting holographic network topologies and transfiguring networks.

6.1 Distributed Filtering

Hyperspeed signaling supports a variety of distributed filtering configurations. The simplest configuration is “egress filtering” that can be used by service provider networks and other entities that transport traffic between networks. As shown in Figure 2, when a malicious packet is identified, a hyperspeed sentinel message is sent to the egress filter to intercept the packet. If the reaction time window is sufficiently large, the sentinel message arrives at the egress filter in advance of the malicious packet to permit the threat to be neutralized. The sentinel message must contain enough information to identify the malicious packet. Note that the malicious traffic is dropped at the egress filter, and the downstream network is unaware of the attempted attack.

Hyperspeed signaling enhances flexibility and efficiency by distributing detection and filtration functionality. Also, it enables service provider networks and other networks (e.g., enterprise networks) that employ multiple detection modalities to maintain low latency. The traditional ingress filtering approach is shown in Figure 3(a). This approach deploys detector-filters in series, where

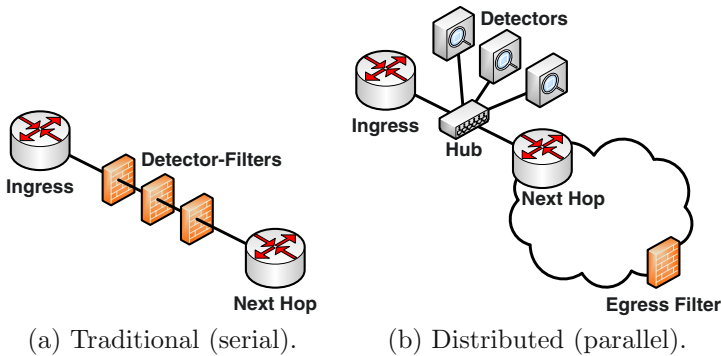


Figure 3. Traditional and distributed filtering configurations.

each detector-filter (e.g., firewall) contributes to the overall delay. On the other hand, the distributed filtering approach shown in Figure 3(b) deploys detectors in parallel at ingress and a filter at egress. Thus, the overall delay is the delay introduced by the single slowest detector plus the delay required for egress filtering.




Figure 4. Advance warning.

Figure 4 shows an advance warning configuration where a hyperspeed signal (sentinel message) is sent to the customer (or peer) ingress instead of the provider egress. In this configuration, the service provider (or peer) network detects malicious packets, but only alerts the customer (or peer) network about the incoming packets. Since the other party has advance warning, it can observe, analyze and/or block the packets or take any other actions it sees fit.

The advance warning configuration enables networks to outsource detection. Copies of suspicious packets could be forwarded to a third party that has sophisticated detection capabilities (e.g., security service providers or government agencies). If the third party detects malicious activity, it can send a hyperspeed signal to trigger filtering. The third party could correlate packets observed from multiple client networks and provide sophisticated detection services to its clients without compromising its intellectual property or national security.




Figure 5. Simple teleportation.

6.2 Teleportation

Hyperspeed routes can be used to teleport packets. Simple teleportation is shown in Figure 5. An operator located at A sends a packet along path ABFG where the hop from B to F involves teleportation. To teleport the packet from B to F, the packet could be encrypted and encapsulated in a labeled ICMP ping packet and sent to B along a hyperspeed path, where it would be converted to its original form and forwarded to G along a normal path. If the teleportation mechanism is to be further concealed, then the packet could be fragmented and the fragments sent along different hyperspeed paths to F (assuming that multiple hyperspeed paths exist from B to F).

Another teleportation approach takes after stage magic. Magicians often use identical twins to create the illusion of teleportation. To set up the act, the magician positions one twin at the source while the other is hidden at the destination. During the act, the magician directs the first twin to enter a box and then secretly signals the other twin to reveal himself at the destination. The same approach can be used to create the illusion of packet teleportation.

The staged teleportation approach is shown in Figure 6. An operator at A uses simple teleportation to secretly send a packet from A to F (Step 1). Next, the operator sends an identical packet from A to B along a normal path; this packet is dropped upon reaching B (Step 2). The operator then sends a hyperspeed signal from A to F (Step 3), which causes the staged packet to move from F to G along a normal path (Step 4). A casual observer would see the packet travel from A to B and the same packet subsequently travel from F to G, but would not see the packet travel from B to F (because no such transmission took place). Depending on the time-sensitivity of the operation, the stage can be set (Step 1) long before the act (Steps 2, 3 and 4) takes place.

A variation of the teleportation act involves a modification of Step 1. An operator located at F sends a copy of a packet to A along a covert hyperspeed

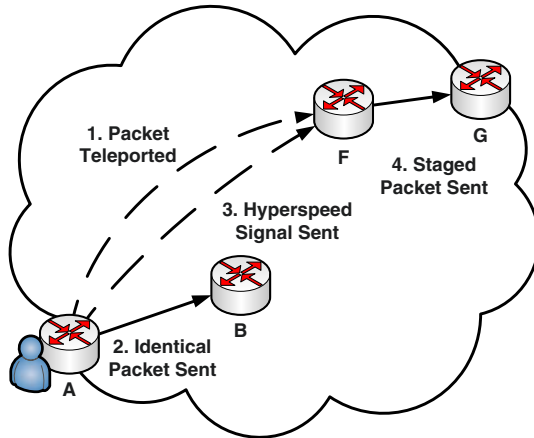


Figure 6. Staged teleportation.

path. As in the previous scenario, a casual observer would see the packet travel from A to B and then from F to G, but not from B to F. This staged teleportation approach can help conceal the real origins of network messages.




Figure 7. Quarantining Network Devices.

6.3 Quarantining

Quarantining enables a network device, segment or path to disappear before it can be compromised by an attack. As shown in Figure 7, a detector located upstream of a targeted device identifies a threat. The detector then sends hyperspeed signals to the appropriate network nodes to prevent malicious traffic from reaching the device. This essentially quarantines the targeted device.

Note that if the attack reaches the targeted device before it is quarantined, the device is isolated before it can affect other parts of the network; the device is reconnected only after it is verified to be secure. Of course, the fact that




Figure 8. Tagging.

quarantine messages travel along hyperspeed paths increases the likelihood that the attack will be thwarted before it impacts the targeted device. The same technique can be used to quarantine network segments or deny the use of certain network paths.

6.4 Tagging

One application of tagging is similar to the use of pheromone trails by animals. In this application, a network essentially tracks the paths taken by suspicious traffic. A network administrator sends diagnostic packets via hyperspeed paths to nodes along the path taken by a suspicious packet in order to observe its behavior. If, as shown in Figure 8, the suspicious packet causes anomalous behavior at one of the nodes, the diagnostic packet reports the anomaly via a hyperspeed signal and the compromised device can be quarantined as described above. In extreme cases, all the nodes on the path taken by the suspicious packet could be quarantined until the situation is resolved.

Tagging can also be used to mitigate the effects of attacks that originate from multiple sources, including distributed denial-of-service attacks (DDoS) and other novel attacks. Consider a sophisticated attack that works like the “five finger death punch” in the movie *Kill Bill Vol. 2*. The attack, which is fragmented into five benign packets, is executed only when all five packets are assembled in sequence. Since a single stateful firewall with knowledge about the fragmented attack could detect and block one or more packets, implementing a successful attack would require the packets to be sent from different locations.

The tagging mechanism can counter the fragmented attack by quarantining the target as soon as anomalous behavior is detected. The packets constituting the attack could then be traced back to their origins at the network perimeter, and security devices configured appropriately to detect the attack.




Figure 9. Network holography.

6.5 Network Holography

Networks can hide their internal structure, for example, by using private IP addresses. Hyperspeed signaling enables networks to project illusory internal structures or “holograms.”

Conventional holograms are created using lasers and special optics to record scenes. In some cases, especially when a cylindrical piece of glass is used, a scene is recorded from many angles. Once recorded, the original scene may be removed, but the hologram still projects the recorded scene according to the viewing angle. If enough angles are recorded, the hologram creates the illusion that the original scene is still in place.

Similarly, an illusory network topology can be created in memory and distributed to edge nodes in a real network (Figure 9). The presence of multiple hyperspeed paths between pairs of edge routers can help simulate the illusory topology. Other nodes may be included, but the edge nodes at the very minimum must be included. When probes (e.g., ping and traceroute) hit the real network, the edge nodes respond to the probes as if the network has the illusory topology. It is important that the same topology is simulated from all angles (i.e., no matter where the probe enters the network) because any inconsistency would shatter the illusion.

6.6 Transfiguration

Transfiguration enables networks to cooperate, much like utilities in the electric power grid [16], to provide services during times of crisis. Network administrators can manipulate their internal network topologies or modify the topologies along the perimeters of cooperating networks to lend or lease additional resources as required. Additionally, administrators may need to modify the topologies at the perimeter boundaries near an attack. This method is analogous to moving the frontline forward or backward during a battle.

Links and nodes may need to be strategically quarantined, disabled or re-enabled based on the circumstances. As resources are lost and gained, the roles of devices, especially at the perimeter, may change. Hyperspeed signaling permits topology changes to occur seemingly instantaneously and enables devices with special roles to operate in proxy where necessary at the perimeter. As resources become available (either regained after being compromised or leased from other sources), the window for hyperspeed signaling can be adjusted to provide additional reaction time.

7. Implementation Challenges

This section discusses the principal challenges involved in implementing hyperspeed signaling in MPLS networks. The challenges include network deployment, traffic burden and net neutrality.

7.1 Network Deployment

Deploying a hyperspeed signaling protocol in a network involves two principal tasks. The first is programming the hardware devices to satisfy the hyperspeed signaling requirements for the specific network. Second, the hardware devices must be installed, configured and calibrated for efficient hyperspeed signaling. Ideally, vendors would program the algorithms/protocols in the hardware devices. The installation, configuration and calibration of the devices would be performed by network engineers and administrators. This task would be simplified and rendered less expensive if vendors were to offer software/firmware updates for deploying hyperspeed signaling without the need to replace existing network devices.

7.2 Traffic Burden

Sending traffic along suboptimal paths essentially increases network “capacity” – keeping more packets in the network at any given time. Because the additional time that a non-hyperspeed packet spends in the network is specified by the reaction time window, the amount of additional traffic flowing in the network is approximately equal to the product of the reaction time window and the average link bandwidth.

Another metric for the burden imposed by hyperspeed signaling is the non-hyperspeed delay divided by the hyperspeed delay (the non-hyperspeed delay is equal to the hyperspeed delay plus the reaction time window). This metric only applies to pairs of communicating end points.

MPLS networks may need additional bandwidth depending on their capacity and the presence of alternate links. The traffic burden due to hyperspeed signaling can be reduced by strategically partitioning a network into multiple signaling domains to prevent critical links from being flooded.

A traffic burden is also induced in a distributed filtering scenario where malicious traffic is allowed to flow through the network and screened later

(e.g., at an interior or egress node). However, this is not an issue because most service provider networks simply transport traffic, leaving the task of filtering to customers.

7.3 Net Neutrality

Issues regarding net neutrality must be considered because the implementation of hyperspeed signaling requires command and control traffic to be treated differently from other traffic. In particular, non-hyperspeed traffic is intentionally slowed to accommodate the desired reaction time windows.

At this time, there is no consensus on the definition of net neutrality [14]. Does net neutrality mean that all traffic should be treated the same? Or does it mean that only traffic associated with a particular application type should be treated the same?

Regardless of its definition, net neutrality is not a major concern for VPN service providers, who can give preferential treatment to traffic based on the applicable service level agreements. In the case of Internet service providers, net neutrality would not be violated as long as all customer traffic is slowed by the same amount.

Currently, no laws have been enacted to enforce net neutrality, although there has been considerable discussion regarding proposed legislation. Many of the proposals permit exceptions in the case of traffic management, public safety and national security. Since hyperspeed signaling, as discussed in this paper, focuses on network defense, it is reasonable to conclude that it would fall under one or more of the three exemptions.

Interestingly, the distributed filtering technique provided by hyperspeed signaling actually enables service providers to treat different types of traffic in a “more neutral” manner than otherwise. Consider a situation where a service provider employs a firewall that performs deep packet inspection. Certain types of traffic (e.g., suspicious packets) would require more inspection time by the firewall, contributing to a larger delay than for other traffic. But this is not the case when all traffic (including suspicious traffic) is allowed to enter the network while copies are simultaneously sent to a distributed detector. Malicious packets are filtered at egress or elsewhere in the network using hyperspeed signaling. Non-malicious packets in the same suspicious traffic pass through the network just like normal traffic.

8. Conclusions

As attacks on computer and telecommunications networks become more prolific and more insidious, it will be increasingly important to deploy novel strategies that give the advantage to network defenders. Hyperspeed signaling is a promising defensive technology that could combat current and future threats. The paradigm is motivated by Arthur C. Clarke’s third law of prediction: “Any sufficiently advanced technology is indistinguishable from magic” [6]. Hyperspeed signaling does not require electrons to move faster than the laws of physics

permit; instead, malicious traffic is slowed down ever so slightly to endow defensive capabilities that are seemingly magical. The hallmark of good engineering is making the right trade-off. Intentionally slowing down network traffic may appear to be counterintuitive, but the defensive advantages gained by hyper-speed signaling may well outweigh the costs.

Note that the views expressed in this paper are those of the authors and do not reflect the official policy or position of the U.S. Department of Defense or the U.S. Government.

References

- [1] L. Anderson, P. Doolan, N. Feldman, A. Fredette and B. Thomas, LDP Specification, RFC 3036, 2001.
- [2] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan and G. Swallow, RSVP-TE: Extensions to RSVP for LSP Tunnels, RFC 3209, 2001.
- [3] T. Bates, Y. Rekhter, R. Chandra and D. Katz, Multiprotocol Extensions for BGP-4, RFC 2858, 2000.
- [4] R. Best, Intelligence, Surveillance and Reconnaissance (ISR) Programs: Issues for Congress, CRS Report for Congress, RL32508, Congressional Research Service, Washington, DC, 2005.
- [5] U. Black, *MPLS and Label Switching Networks*, Prentice Hall, Upper Saddle River, New Jersey, 2002.
- [6] A. Clarke, *Profiles of the Future: An Inquiry into the Limits of the Possible*, Harper and Row, New York, 1999.
- [7] B. Davie and Y. Rekhter, *MPLS: Technology and Applications*, Morgan Kaufmann, San Francisco, California, 2000.
- [8] E. Gray, *MPLS: Implementing the Technology*, Addison-Wesley, Reading, Massachusetts, 2001.
- [9] B. Nicolls, *Airman's Guide*, Stackpole Books, Mechanicsburg, Pennsylvania, 2007.
- [10] L. Peterson and B. Davie, *Computer Networks: A Systems Approach*, Morgan Kaufmann, San Francisco, California, 2003.
- [11] E. Rosen and Y. Rekhter, BGP/MPLS IP Virtual Private Networks (VPNs), RFC 4364, 2006.
- [12] E. Rosen, A. Viswanathan and R. Callon, Multiprotocol Label Switching Architecture, RFC 3031, 2001.
- [13] T. Russell, *Signaling System #7*, McGraw-Hill, New York, 1998.
- [14] H. Travis, The FCC's new theory of the First Amendment, *Santa Clara Law Review*, vol. 51(2), pp. 417–513, 2011.
- [15] United States Department of Defense, Military Deception, Joint Publication 3-13.4, Washington, DC, 2006.
- [16] M. Wald, Hurdles (not financial ones) await electric grid update, *New York Times*, February 6, 2009.