

# Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile

Upul Jayasinghe  
LCA2, I&C  
EPFL, Switzerland  
upul.jayasinghe@epfl.ch

Sergio Barreto  
LCA2, I&C  
EPFL, Switzerland  
sergio.barreto@epfl.ch

Miroslav Popovic  
LCA2, I&C  
EPFL, Switzerland  
miroslav.popovic@epfl.ch

Teklemariam T. Tesfay  
LCA2, I&C  
EPFL, Switzerland  
tech.tesfay@epfl.ch

Jean-Yves Le Boudec  
LCA2, I&C  
EPFL, Switzerland  
jean-  
yves.leboudec@epfl.ch

## ABSTRACT

We are interested in the security of the MPLS Transport Profile (MPLS-TP), in the context of smart-grid communication networks. The security guidelines of the MPLS-TP standards are written in a complex and indirect way, which led us to pose as hypothesis that vendor solutions might not implement them satisfactorily. To test this hypothesis, we investigated the Cisco implementation of two MPLS-TP OAM (Operations, Administration, and Maintenance) protocols: bidirectional forwarding detection (BFD), used to detect failures in label-switched paths (LSPs) and protection state coordination (PSC), used to coordinate protection switching. Critical smart grid applications, such as protection and control, rely on the protection switching feature controlled by BFD and PSC. We did find security issues with this implementation. We implemented a testbed with eight nodes that run the MPLS-TP enabled Cisco IOS; we demonstrated that an attacker who has access to only one cable (for two attacks) or two cables (for one attack) is able to harm the network at several points (e.g., disabling both working and protection LSPs). This occurred in spite of us implementing the security guidelines that are available from Cisco for IOS and MPLS-TP. The attacks use forged BFD or PSC messages, which induce a label-edge router (LER) into believing false information about an LSP. In one attack, the LER disables the operational LSP; in another attack, the LER continues to believe that a physically destroyed LSP is up and running; in yet another attack, both operational and backup LSPs are brought down. Our findings suggest that the MPLS-TP standard should be more explicit when it comes to security. For example, to thwart the attacks revealed here, it should mandate either hop by hop authentication (such as MACSec) at every node, or an ad-hoc authentication mechanism for BFD and PSC.

## Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Miscellaneous

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).  
*SEGS'14*, November 7, 2014, Scottsdale, Arizona, USA.  
Copyright 2014 ACM 978-1-4503-2954-8/14/11 ...\$15.00.  
<http://dx.doi.org/10.1145/2667190.2667197>.

## Keywords

MPLS-TP; CISCO IOS; Smart Grid Security; Vulnerability; Authentication

## 1. INTRODUCTION

The MPLS Transport Profile (MPLS-TP) is one of the proposed communication technologies for smart-grid networks [6]. MPLS-TP is mainly used for long-distance inter-control center communication of measurement data (e.g. synchrophasor data from Phasor Measurement Units), and control and protection commands. Given the critical nature of these type of communications for the reliable operation of a smart grid, the communication infrastructure is required to satisfy high availability with bounded delay. MPLS-TP satisfies these requirements in that it supports traffic engineering to guarantee deterministic delay for high priority traffic, and it provides end-to-end protection - ensuring network reliability and high availability. End-to-end protection is achieved by the MPLS-TP OAM (Operations, Administration and Maintenance) framework that provides protection switching feature, controlled by bidirectional forwarding detection (BFD) and protection state coordination (PSC) protocols. BFD detects failures in label-switched paths (LSP), and PSC coordinates the protection switching.

The fact that an MPLS-TP network extends over a (usually unprotected) wide area renders the communication network vulnerable to cyber intrusions by an attacker with a malicious intention of compromising the smart grid's operations. Hence, one of the major challenges for a smart grid utility is to implement proper cyber security protection methods for its MPLS-TP based WAN.

There is a whole family of MPLS-TP-related RFCs [7, 8, 12, 15, 16] and several among them are security-related. However, we find that the RFC-based security analysis of MPLS-TP is complex due to fragmentation of pieces of information needed to understand the big picture of which the pieces are spread among multiple RFCs. To some extent, security-related problems are treated as "hot potatoes"; the responsibility of securing different aspects of MPLS-TP in different RFCs is sometimes outsourced to another RFC some of which are only informational and majority of which is without straightforward guidelines. Consequently, the lack of holistic approach in securing the MPLS-TP network makes the whole process difficult to follow. For example, RFC 5085 [12] relies on IPsec to provide the security of MPLS-TP, but this is not always applicable as there are some non IP flows in some smart grid and other contexts that use MPLS-TP. This led us to pose as a hypothesis that

vendor solutions might not correctly implement all of the required security. To test this hypothesis, we started an analysis of several MPLS-TP implementations. In this paper, we report on our findings with the Cisco IOS implementation, Cisco being one of the leading manufacturers of network equipment and with a large investment in smart grids.

In Section 2, we give an overview of the MPLS-TP features that are necessary to understand the experiments we conduct. In Section 3 we describe the testbed, that we build in the lab environment: it consists of eight virtual routers that each runs a Cisco IOS image that supports MPLS-TP.

In Section 4, we describe how we were able to conduct the attacks. This occurred in spite of us implementing the security guidelines that are available from Cisco for IOS and MPLS-TP (which include IPSec and authentication of the control plane). In some scenarios (not reported here) the available security measures were sufficient, however, we identified other scenarios, described below, where these security measures were not sufficient. In these latter scenarios, the attacker harms the network at several points (e.g., disabling both working and protection LSPs) by using the access to only one or two interfaces of a switch (e.g., on the protection LSP). We achieved this by inserting forged BFD or PSC messages into the network, which induces the label edge router (LER) into believing false information about the LSP status. In two attacks, the LER disables the operational LSP. In another attack, the LER continues to believe that a physically destroyed LSP is up and running. These spoofing attacks rely on the assumption that an attacker gains physical access to cables in the network. This assumption is consistent with reality because it is common to find several unmanned facilities, such as remote substations, in smart grid networks. Such substations are where fiber-to-copper converters or optical terminators of an MPLS-TP network likely reside. Moreover, some smart grid utilities also deploy pole-mounted optical repeaters at every few kilometres along their electric transmission lines. Gaining physical access to such physically exposed locations is usually as easy as breaking a window in a substation or climbing up a pole.

In addition to the ease of physical access to the attack locations, the equipment required to mount the spoofing attacks is rather affordable. In MPLS-TP networks where only optical links are used (no copper) an attacker can afford a NetFPGA-10G card [10] with SFP+ modules in order to gain access to the network. Besides, it is not uncommon to find copper cables between optical termination nodes (fiber-to-copper converters) and MPLS-TP routers (see Figure 4). The cost of equipment required to launch a spoofing attack in such networks is almost negligible. An attacker can utilize an off-the-shelf switch to connect his laptop computer to the network at the copper cable segments. The attacker only injects very modest amounts of bogus traffic from his laptop computer and does not need to intercept legitimate traffic, which remains untouched. As such, our described attacks require only low-end, cheap equipment that is readily available in all consumer electronics shops. In our experiments, we used the second method to access the network in order to demonstrate the spoofing attacks discussed in this paper. Note that the nature of attacks remains the same no matter which method of gaining access to the network is used.

In two of our attacks, the attacker needs to access only one cable; in another one, he needs access to two cables. In all cases, an attacker gains more power to damage the physical electrical infrastructure as a result of a more intelligent communication network, i.e., an attacker can now bring down a physical target in the power grid from a remote location by selectively manipulating traffic at only one, or in some cases two, locations in the communication infrastructure. In conventional power grid networks, an attacker

would have required physical access to sabotage the target. If we want to compare the attacks described in this paper to the one that involves simple wire cutting, we can see that cutting wires harms only LSPs directly affected by the cut whereas we show that our attacks are more powerful (LSPs in the other parts of the network are also affected). In some sense, this is in contradiction with the expected benefits of smart grids: they should not make the electrical systems weaker than they are today. In Section 5 we discuss countermeasures that are likely to thwart the described attacks.

## 2. MPLS-TP PROTOCOL OVERVIEW

A label-switched path (LSP) is defined as a one-way path that data follows from one particular node (a router able to do label-switching) to a different node, where intermediate nodes can be traversed. The two nodes where data enters and leaves the LSP are called label edge routers (LER), and nodes traversed within the LSP are called label-switching routers (LSR). MPLS-TP mandates that all LSPs go by pairs traversing the same nodes and links on each direction, and that they be signaled as a single entity (one single LSP identifier is used); this characteristic is called co-routed bidirectional.

In MPLS-TP, the control plane for signalling and recovering LSPs can be static or dynamically configured. In smart grids, a static configuration is commonly used by network operators for security reasons, to avoid interacting with dynamic control protocols (i.e., RSVP-TE, T-LDP) from other service providers; as we are in the context of smart-grid communication networks, we use static configuration of MPLS-TP LSPs for our analysis.

Recovery is the ability of the network to become operational following the failure or degradation of traffic delivery caused by a network fault or a denial-of-service attack on the network [15]. There are several types of recovery methods in MPLS-TP, and in our testbed we analyze protection switching. Protection switching is a well-suited method that provides fast repair and exists side-by-side with the static configuration of the control plane commonly used in smart grids. This method pre-allocates an alternative bidirectional LSP to divert traffic during a fault condition, uses BFD for link failure detection with an almost immediate response time (i.e., less than 50ms) and has a mechanism for coordinating the state of the protection provided to a working LSP between both LERs. Within the protection switching scope, a working LSP is defined as the LSP where data runs under normal operating conditions; furthermore, a protection LSP is the pre-allocated LSP where data is diverted from the working LSP in case of network failure or degradation. The service delivered by both working and protection LSPs is called an "MPLS-TP tunnel".

There are two different schemes for providing protection switching to an LSP. The first scheme is called "1+1", where the label edge router at the ingress of the MPLS-TP tunnel transmits simultaneously the data on both working and protection LSPs, and the label edge router at the egress of the MPLS-TP tunnel selects between working or protection LSPs based on some predetermined criteria. The second scheme is called "1:n" where the working LSP handles data under normal operating conditions; and only if there is a defect, failure, degradation or request from network operator, the traffic is switched to protection LSP. For simplicity, in this paper we discuss a particular case of the second scheme called "1:1" where one pre-allocated protection LSP serves one particular working LSP. The nature of the attacks we conduct is such that the security weaknesses revealed in our testbed can be exploited in schemes "1+1" and "1:n" as well, with exactly the same results.

Ver	Diag	Sta	P	F	C	A	D	M	Detect Mult	Length
My Discriminator										
Your Discriminator										
Desired Min TX Interval										
Required Min RX Interval										
Required Min Echo RX Interval										
Auth Type		Auth Length		Authentication Data						

Figure 1: Format of a BFD Control Packet.

## 2.1 Bidirectional Forwarding Detection (BFD)

BFD in MPLS-TP is a protocol that detects link failures within the MPLS-TP tunnel. From the three types of messages described in [1], our focus is on the BFD control packet that we are spoofing in our attack. The BFD control packet verifies the continuity of an LSP. A BFD session is established between label edge routers for each LSP within the MPLS-TP tunnel. A LER sends a BFD control packet every 3.3 ms (this is the interval recommended by IETF [15]); the BFD control packet is sent in-band in the LSP (i.e. is switched at intermediate routers exactly like data packets in the LSP) and is intercepted by the LER that terminates the LSP. When an LER observes that 3 consecutive BFD control packets are not received, it declares the incoming LSP to be broken; as the two directions of co-routed bidirectional LSPs can fail independently, the receiving LER sends a BFD “Session Down” message in the reverse direction of the LSP to inform the LER at the other end of the failure; then protection switching is triggered. In Cisco’s implementation BFD control packets are sent every 4 ms instead of 3.3 ms.

In order to understand the attacks described in section 4, we include the format of a BFD control packet in Figure 1 and provide relevant information for the fields concerning our attacks. The first of two fields for our testbed is the diagnostics (*Diag*) field, which is five bits long and reports a fault or defect condition between label edge routers; from the 32 possible codes we are interested in two of them, a 0 means there is no fault or defect condition to report and a code 1 stands for “Control Detection Timer Expired”, which in normal operating conditions is set when we miss three consecutive BFD control packets. The second field is the state (*Sta*) field, which is two bits long and refers to the state of the BFD session between label edge routers; here a value of zero means “Administrative Down” (given by network operator), one stands for “Down”, two for “Init” (used during BFD session setup) and three for “Up”.

## 2.2 Protection State Coordination (PSC)

The PSC protocol is used to ensure that — whenever protection switching is triggered in one of the unidirectional LSPs of an MPLS-TP tunnel — the protection switching is also triggered for the remaining unidirectional LSP. This is in line with the co-routed bidirectional feature of MPLS-TP LSPs described in Section 2.1. PSC protocol also tells the LER whether the protection LSP is available, and if there is any inconsistency in protection switching configuration (timers, revertive functionality, etc.) between LERs.

There are six different PSC protocol states, among which the normal state is the default state when protection switching is enabled; and the unavailable state, which is used when protection switching is disabled by network operator or unavailable due to a failure on the protection LSP. A label edge router calculates the next PSC protocol state, based on the priorities of the requests issued by three sources: local requests (which can come from the network operator, control plane, management plane or specific timers), the PSC message received from the peer label edge router, and the

Ver	Request	Prot. Type	Revertive	Reserved1	FPath	Path
TLV Length			Reserved2			
Optional TLVs						

Figure 2: Format of a PSC Control Packet.

current PSC protocol state. According to [16], the highest priority of the requests issued by any of the sources described above corresponds to network-operator commands (“Clear”, “Lockout of Protection” and “Forced Switch”). Network-operator commands are used by the attackers to launch the attacks on the PSC protocol described in Section 4.2. For these attacks, we discuss two fields of the PSC control packet format (Figure 2). First we have the 4-bit request (*Request*) field that represents the PSC protocol state of the local label edge router that is sent to the peer label edge router to be considered on its next state computation. In this field code 14 stands for “Lockout of Protection” and indicates that protection switching is down as a result of a network operator command. The second field is the 8-bit fault path (*FPath*) field that indicates which LSP (working or protection) shall be affected by the *Request* field.

## 3. TESTBED DESCRIPTION

In this section, we describe the setting we used to evaluate MPLS-TP security. The network topology that we used is shown in Figure 3 [4]. Depicted routers run Cisco IOS that supports MPLS-TP; namely, we use Cisco Cloud Service router 1000V (CSR1000V) IOS. We mount eight VMWare virtual machines (VM) configured with CSR 1000V images on two physical machines (four virtual routers per each of the two physical machines). Each physical machine has an eight-core processor and 16GB of RAM. We assign one processor core and 3GB of RAM for each of the four virtual routers within one physical machine. The CSR 1000V 60-day evaluation license that we had at our disposal gave us full access to all the CSR 1000V features at a throughput of 50 Mbps. To the best of our knowledge, Cisco’s other commercially available routers that support MPLS-TP have no additional security-related features.

We configure our MPLS-TP network to follow a one-to-one (1:1) protection mechanism by configuring a working LSP and a protection LSP between R2 and R7. A working LSP follows the path R2-R1-R3-R5-R7, whereas a protection LSP follows the path R2-R4-R6-R8-R7. We configure LSPs statically and use RSVP to reserve network resources. In order to connect each virtual router, as well as two physical machines, we use 1 Gbps links with full duplex configuration. As MPLS TP recommends co-routed bidirectional LSPs as described in Section 2, both directions of each link are assigned 25Mbps of bandwidth. Inside both the working and protection LSPs, we configure a BFD session, described in Section 2.1, with a 4ms message interval and a 12ms detection interval [5].




Figure 3: The Network topology with 1:1 Protection used in our MPLS-TP testbed.




Figure 4: BFD Spoofing Attack : Removing protection from a target LSP.

## 4. VULNERABILITIES IN MPLS-TP PROTOCOL

In this section, we describe three attacks. In Section 4.1 we describe two BFD spoofing attacks: In the first one, we remove protection from a target LSP; and in the second one, we disable fault detection in an LSP. In Section 4.2 we describe a PSC spoofing attack in which we bring down an operational MPLS-TP tunnel. As discussed in Section 1, we assume that the attacker can access the cables at one point (for the first and third attacks) or two points (for the second attack).

### 4.1 BFD Spoofing Attacks

As described in Section 2.1, BFD control messages are used to proactively monitor the continuity of an LSP. In this section, we describe spoofing attacks associated with BFD messages that enable an attacker to launch targeted attacks on a specific LSP.

#### 4.1.1 Scenario I - Removing Protection from a Target LSP

In this attack, the attacker’s goal is to falsely inform a label edge router that a working LSP is broken, even though there is no actual failure. This attack forces label edge routers of an LSP to unnecessarily switch from a working LSP to a protection LSP.

In our experiment, first we connected our switch to the cable of the working LSP at point *a* in Figure 4. Then we connected our laptop *b* to the switch and sniffed for BFD packets in order to gather information such as MAC addresses of LSRs *R3* and *R5*, MPLS label, the TTL value, and the BFD session number of the target LSP between LERs *R2* and *R7*. Finally, using the sniffed information, we created the forged packets by using the Scapy [13] tool and sent them to *R7* through the switch by using a simple python code. The packets were manipulated such that we modified the *Diagnostic (Diag)* field and the *State (Sta)* field of a BFD control packet shown in Figure 1; the *Diag* field was set to “Control Detection Time Expired” and the *Sta* to “Down” (Figure 5). Note that other *Diag* field codes could also be used with similar results.




Figure 5: A wireshark capture of a spoofed BFD packet to remove protection from a target LSP.

Upon reception of *BFD Down* packets, *R7* is deceived into believing that the forward working LSP from *R2* to *R7* is down. In order to observe the result of this attack, we sniffed at points *c* and *d* of Figure 4<sup>1</sup>; we observed that *R7* starts the protection switching by sending to *R2* three PSC packets via the protection LSP with a “Signal Fail (*SF*)” message, and by sending BFD control packets via the working LSP with the *Sta* field set to “Down”. *R7* does this despite the fact that it also receives legitimate BFD packets from *R2* with *Sta* field set to “Up”. When *R2* receives the BFD and PSC messages from *R7* through the working LSP and protection LSP (respectively), it triggers the protection switching and sends to *R7* three PSC control packets with a “Signal Fail (*SF*)” message.

We observed that, by continuously injecting the forged BFD Down packets destined to *R7*, we impede the working LSP from getting back on feet - effectively removing protection from the target LSP. This is possible because there is no mechanism for detecting that the frequency with which the messages are coming is different from what is expected. In other words, no matter how many forged messages we insert within the expected period of *4ms*, it does not raise any suspicion and those messages are accepted as legitimate. Similarly, no suspicion is raised when a node receives a steady mixture of “Up” and “Down” messages. As a final outcome of the attack, in the case of a fault in the protection LSP, the LERs would not have any alternative LSP to switch to.

#### 4.1.2 Scenario II - Disabling LSP Fault Detection

The attacker’s goal in this attack is the inverse of the attack introduced above, i.e., he aims to falsely inform label edge routers that a working LSP is up and running, while in reality it is down due to a link failure somewhere along the path. The link failure can be due to deliberate sabotage or a result of a random failure. In our experiment, we used the set up in Figure 6 to demonstrate this attack. On the working LSP, we connected two malicious switches *b* and *d* to the cables between *R1* and *R2* and between *R5* and *R7*. Then we connected two laptops *a* and *c* to the MPLS-TP network through these two switches.




Figure 6: BFD Spoofing Attack: Disabling fault detection in an LSP.

From laptop *a*, we injected forged BFD “*Session Up*” packets (the *Sta* field set to “Up”) destined to LER *R2* as if they were sent from LER *R7*. Likewise, we injected forged “*Session Up*” packets from laptop *c* destined to LER *R7* as if they were sent from LER *R2*. The two LERs processed these forged packets without raising any alarms in spite of receiving more BFD “*Session Up*” packets than expected, as a result of the packets injected from our laptops.

We then silently broke the link between *R1* and *R3*, thus emulating a random link failure condition, while we continued sending the forged “*Session Up*” BFD packets to both label edge routers *R2* and *R7* from our two laptops. Again, we observed that the two label edge routers *R2* and *R7* did not notice the change in the rate

<sup>1</sup>Note that packet sniffing at points *c* and *d* is not part of the attack, we used it only to establish that the attack works.



If an MPLS-TP core network supports IP services, OAM messages such as BFD and PSC messages can be tunnelled on top of an IP tunnel. In such cases, standard IP security solutions such as IPsec or (D)TLS between label edge routers can be used to provide end-to-end security, thereby preventing spoofing attacks. RFC5085 [12] proposes IPsec as a solution to protect OAM protocols of MPLS/GMPLS networks if the core network supports IP, VPN, or transport services. However, not all core networks are required to support IP. For example, a smart grid MPLS-TP network that transports IEC 61850 based Multicast Sampled Value (MSV) and Generic Object Oriented Substation Event (GOOSE) messages between substations or between a substation and a controller is often implemented without IP [14].

An alternative solution for preventing spoofing attacks on BFD and PSC messages in non-IP MPLS core networks is to use hop-by-hop security (MACsec). However, Cisco does not support MACsec in its routers. One minor drawback of using MACsec is that if any of the network devices in an LSP are compromised, MACsec fails to achieve its purpose, i.e., forged BFD and/or PSC packets injected at the compromised device will be processed as valid packets by a receiving label edge router. Nonetheless, such attacks can be prevented by incorporating tamper resistant security solutions such as Trusted Platform Module (TPM) to protect sensitive data and by enforcing proper access control mechanisms to deny unauthorised access to the network devices. Note that implementing an ACL alone would not solve the problem as the attack is conducted with spoofed packets. Since our findings show that lack of MAC layer security exposes smart grid networks to various cyber-attacks, we recommend that utilities implement MACsec or a variant of it in their MPLS-TP networks.

## 6. CONCLUSION

MPLS-TP is the proposed technology for WAN connectivity in the context of smart grid. In this paper, we have shown that, when it comes to the security aspects of the standard, there is a discrepancy between RFCs and the Cisco IOS implementation for MPLS-TP we evaluated, which is not surprising given the complexity of RFCs. More specifically, we have observed that the Cisco IOS does not implement security recommendations for OAM protocols, such as BFD and PSC, thus exposing them to different spoofing attacks. In our testbed, to launch spoofing attacks on these two OAM protocols we exploited the identified security vulnerabilities. By launching spoofing attacks, we have shown that we could degrade the performance of an MPLS-TP network by removing a protection LSP of an MPLS-TP tunnel. We have also demonstrated that we can disable detection of a link failure in LSP by tricking label edge routers into believing a failed link is still up and running. Finally we have shown that we can bring the whole MPLS-TP tunnel down by sending forged operator PSC commands.

Our experiments with the Cisco IOS show that no protection against spoofing attacks is provided for non-IP MPLS-TP OAM messages. Therefore, we recommend that RFCs be more directive in proposing built-in security for OAM protocols. They should mandate source authentication mechanisms for both BFD and PSC messages or mandate MACsec or a variant of it as an alternative authentication solution when built-in security is absent.

## 7. REFERENCES

- [1] D. Allan, G. S. Ed., and J. D. Ed. Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile. IETF, RFC 6428 (Proposed Standard), Nov. 2011. Updated by RFC 7214.
- [2] L. Andersson, L. Berger, L. Fang, N. Bitar, and E. Gray. MPLS Transport Profile (MPLS-TP) Control Plane Framework. IETF, RFC 6373 (Informational), Sept. 2011.
- [3] Cisco Systems, Inc. Cisco SAFE Reference Guide. [http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/SAFE\\_RG/SAFE\\_rg.html](http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html), Jul. 2010.
- [4] Cisco Systems, Inc. Cisco CSR 1000V Series Cloud Services Router Overview. <http://www.cisco.com/c/en/us/td/docs/routers/csr1000/software/configuration/csr1000Vswcfg/csroverview.html>, Jul. 2012.
- [5] Cisco Systems, Inc. IP Routing BFD Configuration Guide, Cisco IOS XE Release 3S. [http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute\\_bfd/configuration/xe-3s/irb-xe-3s-book.html](http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-3s/irb-xe-3s-book.html), 2013.
- [6] A. Durai and V. Varakantam. Building Smart Grid Core Networks. IEEE Smart Grid Newsletter, October 2012.
- [7] L. Fang. Security Framework for MPLS and GMPLS Networks. IETF, RFC 5920 (Informational), July 2010.
- [8] L. Fang, B. Niven-Jenkins, S. Mansfield, and R. Graveman. MPLS Transport Profile (MPLS-TP) Security Framework. IETF, RFC 6941 (Informational), Apr. 2013.
- [9] D. Katz and D. Ward. Bidirectional Forwarding Detection (BFD). IETF, RFC 5880 (Proposed Standard), June 2010.
- [10] J. W. Lockwood, N. McKeown, G. Watson, G. Gibb, P. Hartke, J. Naous, R. Raghuraman, and J. Luo. NetFPGA—An Open Platform for Gigabit-Rate Network Switching and Routing. In *Microelectronic Systems Education, 2007. MSE '07. IEEE International Conference on*, pages 160–161, June 2007.
- [11] S. Mansfield, E. Gray, and K. Lam. Network Management Framework for MPLS-based Transport Networks. IETF, RFC 5950 (Informational), Sept. 2010.
- [12] T. Nadeau and C. Pignataro. Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires. IETF, RFC 5085 (Proposed Standard), Dec. 2007. Updated by RFC 5586.
- [13] Philippe Biondi. Scapy Project. <http://www.secdev.org/projects/scapy/>, Feb. 2011.
- [14] IEC TC/SC 57. IEC 61850 Communication networks and systems for power utility automation - Part 8-1: Specific communication service mapping (SCSM) - Mappings to MMS and to ISO/IEC 8802-3, Jun. 2011.
- [15] N. Sprecher and A. Farrel. MPLS Transport Profile (MPLS-TP) Survivability Framework. IETF, RFC 6372 (Informational), Sept. 2011.
- [16] Y. Weingarten, S. Bryant, E. Osborne, N. Sprecher, and A. Fulignoli. MPLS Transport Profile (MPLS-TP) Linear Protection. IETF, RFC 6378 (Proposed Standard), Oct. 2011. Updated by RFCs 7214, 7271, 7324.