

# The Internet Protocol Journal

March 2013

Volume 16, Number 1

*A Quarterly Technical Publication for  
Internet and Intranet Professionals*

## In This Issue

From the Editor .....	1
SDN and OpenFlow .....	2
Address Authentication .....	15
WCIT Report .....	21
Letters to the Editor.....	34
Book Review.....	36
Fragments .....	38
Call for Papers.....	39

## FROM THE EDITOR

This is the 60th edition of *The Internet Protocol Journal*, and in June we will celebrate our 15th anniversary. Fifteen years is not a long time in absolute terms, but when it comes to networking technology a lot can happen in a short time.

Throughout this 15-year period we have published numerous articles on “emerging technologies,” and in this issue we present yet another. *Software-Defined Networks* (SDNs) have become a mainstream topic for research, development, and standardization. We asked William Stallings to give us an overview of SDNs, and we plan further articles on this topic in the future.

A recurring theme in this journal has been Internet *security* at all levels of the protocol stack. We have covered security in routing, securing the *Domain Name System* (DNS), secure wireless networks, secure HTTP, and much more. This time, Scott Hogg discusses the advantages and disadvantages of using IPv4 or IPv6 addresses as a form of user authentication.

In our previous issue we published some reactions to the outcomes of the *World Conference on International Telecommunications* (WCIT) held in Dubai in December 2012. In this edition, Robert Pepper and Chip Sharp provide analysis and background on this conference and discuss how the revised *International Telecommunication Regulations* (ITRs) might affect the future of the Internet.

It has been some time since we have published a book review, but we are happy to bring you one in this issue. For the first time in history, we are reviewing a book that exists only in electronic form, another sign of a rapidly changing technology landscape. We are always looking for new book reviews. Please send your reviews, letters to the editor, or any subscription questions to [ipj@cisco.com](mailto:ipj@cisco.com)

If you want to look back at 15 years of IPJ, visit our website at [www.cisco.com/ipj](http://www.cisco.com/ipj) where you will find all of our back issues (as a single PDF file, as a collection of individual PDF files, or in HTML format), as well as an index of all IPJ articles.

—Ole J. Jacobsen, Editor and Publisher  
[ole@cisco.com](mailto:ole@cisco.com)

You can download IPJ  
back issues and find  
subscription information at:  
[www.cisco.com/ipj](http://www.cisco.com/ipj)

ISSN 1944-1134

# Software-Defined Networks and OpenFlow

by William Stallings

A network organizing technique that has come to recent prominence is the *Software-Defined Network* (SDN)<sup>[1]</sup>. In essence, an SDN separates the data and control functions of networking devices, such as routers, packet switches, and LAN switches, with a well-defined *Application Programming Interface* (API) between the two. In contrast, in most large enterprise networks, routers and other network devices encompass both data and control functions, making it difficult to adjust the network infrastructure and operation to large-scale addition of end systems, virtual machines, and virtual networks. In this article we examine the characteristics of an SDN, and then describe the *OpenFlow* specification, which is becoming the standard way of implementing an SDN.

## Evolving Network Requirements

Before looking in more detail at SDNs, let us examine the evolving network requirements that lead to a demand for a flexible, response approach to controlling traffic flows within a network or the Internet.

One key leading factor is the increasingly widespread use of *Server Virtualization*. In essence, server virtualization masks server resources, including the number and identity of individual physical servers, processors, and operating systems, from server users. This masking makes it possible to partition a single machine into multiple, independent servers, conserving hardware resources. It also makes it possible to migrate a server quickly from one machine to another for load balancing or for dynamic switchover in the case of machine failure. Server virtualization has become a central element in dealing with “big data” applications and in implementing cloud computing infrastructures. But it creates problems with traditional network architectures (for example, refer to [2]). One problem is configuring *Virtual LANs* (VLANs). Network managers need to make sure the VLAN used by the *Virtual Machine* is assigned to the same switch port as the physical server running the virtual machine. But with the virtual machine being movable, it is necessary to reconfigure the VLAN every time that a virtual server is moved. In general terms, to match the flexibility of server virtualization, the network manager needs to be able to dynamically add, drop, and change network resources and profiles. This process is difficult to do with conventional network switches, in which the control logic for each switch is co-located with the switching logic.

Another effect of server virtualization is that traffic flows differ substantially from the traditional client-server model. Typically, there is a considerable amount of traffic among virtual servers, for such purposes as maintaining consistent images of the database and invoking security functions such as access control. These server-to-server flows change in location and intensity over time, demanding a flexible approach to managing network resources.

Another factor leading to the need for rapid response in allocating network resources is the increasing use by employees of mobile devices such as smartphones, tablets, and notebooks to access enterprise resources. Network managers must be able to respond to rapidly changing resource, *Quality of Service* (QoS), and security requirements.


Existing network infrastructures can respond to changing requirements for the management of traffic flows, providing differentiated QoS levels and security levels for individual flows, but the process can be very time-consuming if the enterprise network is large and/or involves network devices from multiple vendors. The network manager must configure each vendor's equipment separately, and adjust performance and security parameters on a per-session, per-application basis. In a large enterprise, every time a new virtual machine is brought up, it can take hours or even days for network managers to do the necessary reconfiguration<sup>[3]</sup>.

This state of affairs has been compared to the mainframe era of computing<sup>[4]</sup>. In the era of the mainframe, applications, the operating system, and the hardware were vertically integrated and provided by a single vendor. All of these ingredients were proprietary and closed, leading to slow innovation. Today, most computer platforms use the x86 instruction set, and a variety of operating systems (Windows, Linux, or Mac OS) run on top of the hardware. The OS provides APIs that enable outside providers to develop applications, leading to rapid innovation and deployment. In a similar fashion, commercial networking devices have proprietary features and specialized control planes and hardware, all vertically integrated on the switch. As will be seen, the SDN architecture and the OpenFlow standard provide an open architecture in which control functions are separated from the network device and placed in accessible control servers. This setup enables the underlying infrastructure to be abstracted for applications and network services, enabling the network to be treated as a logical entity.

### **SDN Architecture**

Figure 1 illustrates the logical structure of an SDN. A central controller performs all complex functions, including routing, naming, policy declaration, and security checks. This plane constitutes the *SDN Control Plane*, and consists of one or more SDN servers.

Figure 1: SDN Logical Structure



The *SDN Controller* defines the data flows that occur in the *SDN Data Plane*. Each flow through the network must first get permission from the controller, which verifies that the communication is permissible by the network policy. If the controller allows a flow, it computes a route for the flow to take, and adds an entry for that flow in each of the switches along the path. With all complex functions subsumed by the controller, switches simply manage flow tables whose entries can be populated only by the controller. Communication between the controller and the switches uses a standardized protocol and API. Most commonly this interface is the OpenFlow specification, discussed subsequently.

The SDN architecture is remarkably flexible; it can operate with different types of switches and at different protocol layers. SDN controllers and switches can be implemented for Ethernet switches (Layer 2), Internet routers (Layer 3), transport (Layer 4) switching, or application layer switching and routing. SDN relies on the common functions found on networking devices, which essentially involve forwarding packets based on some form of flow definition.

In an SDN architecture, a switch performs the following functions:

- The switch encapsulates and forwards the first packet of a flow to an SDN controller, enabling the controller to decide whether the flow should be added to the switch flow table.
- The switch forwards incoming packets out the appropriate port based on the flow table. The flow table may include priority information dictated by the controller.
- The switch can drop packets on a particular flow, temporarily or permanently, as dictated by the controller. Packet dropping can be used for security purposes, curbing *Denial-of-Service* (DoS) attacks or traffic management requirements.

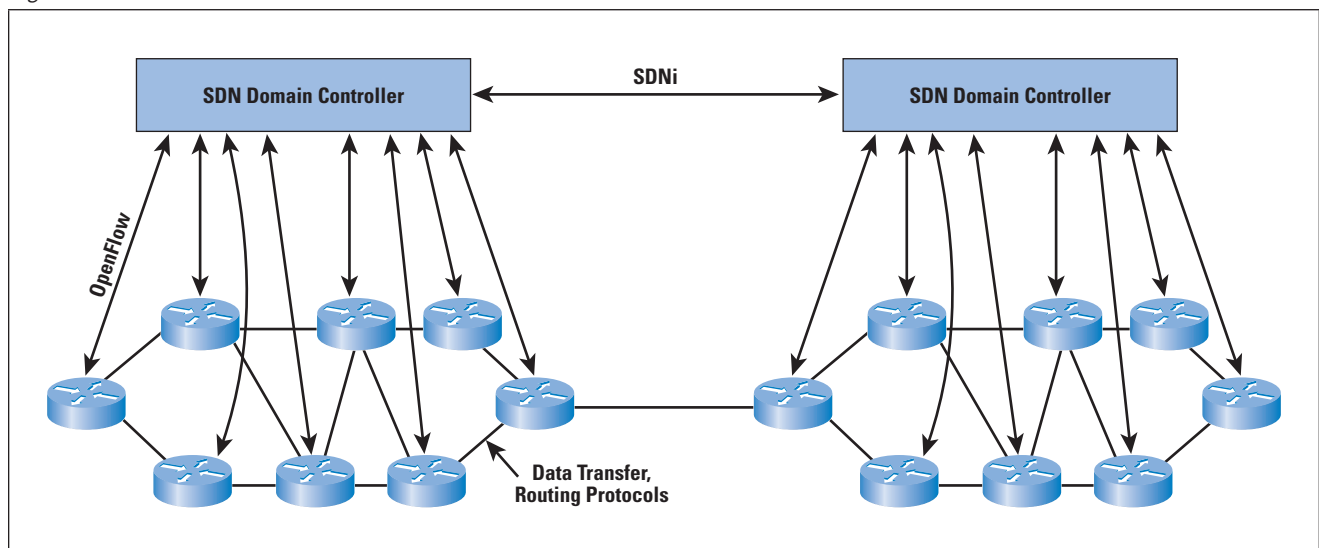
In simple terms, the SDN controller manages the forwarding state of the switches in the SDN. This management is done through a vendor-neutral API that allows the controller to address a wide variety of operator requirements without changing any of the lower-level aspects of the network, including topology.

With the decoupling of the control and data planes, SDN enables applications to deal with a single abstracted network device without concern for the details of how the device operates. Network applications see a single API to the controller. Thus it is possible to quickly create and deploy new applications to orchestrate network traffic flow to meet specific enterprise requirements for performance or security.

### SDN Domains

In a large enterprise network, the deployment of a single controller to manage all network devices would prove unwieldy or undesirable. A more likely scenario is that the operator of a large enterprise or carrier network divides the whole network into numerous nonoverlapping SDN domains as shown in Figure 2.

Figure 2: SDN Domain Structure



Reasons for using SDN domains include the following:

- *Scalability*: The number of devices an SDN controller can feasibly manage is limited. Thus, a reasonably large network may need to deploy multiple SDN controllers.
- *Privacy*: A carrier may choose to implement different privacy policies in different SDN domains. For example, an SDN domain may be dedicated to a set of customers who implement their own highly customized privacy policies, requiring that some networking information in this domain (for example, network topology) not be disclosed to an external entity.
- *Incremental deployment*: A carrier's network may consist of portions of traditional and newer infrastructure. Dividing the network into multiple, individually manageable SDN domains allows for flexible incremental deployment.

The existence of multiple domains creates a requirement for individual controllers to communicate with each other via a standardized protocol to exchange routing information. The IETF is currently working on developing a protocol, called *SDNi*, for “interfacing SDN Domain Controllers”<sup>[5]</sup>. SDNi functions include:

- Coordinate flow setup originated by applications containing information such as path requirement, QoS, and service-level agreements across multiple SDN domains.
- Exchange reachability information to facilitate inter-SDN routing. This information exchange will allow a single flow to traverse multiple SDNs and have each controller select the most appropriate path when multiple such paths are available.

The message types for SDNi tentatively include the following:

- Reachability update
- Flow setup/tear-down/update request (including application capability requirements such as QoS, data rate, latency etc.)
- Capability update (including network-related capabilities such as data rate and QoS, and system and software capabilities available inside the domain)

### OpenFlow

To turn the concept of SDN into practical implementation, two requirements must be met. First, there must be a common logical architecture in all switches, routers, and other network devices to be managed by an SDN controller. This logical architecture may be implemented in different ways on different vendor equipment and in different types of network devices, so long as the SDN controller sees a uniform logical switch function. Second, a standard, secure protocol is needed between the SDN controller and the network device.


Both of these requirements are addressed by *OpenFlow*, which is both a protocol between SDN controllers and network devices, as well as a specification of the logical structure of the network switch functions<sup>[6, 7]</sup>. OpenFlow is defined in the *OpenFlow Switch Specification*, published by the *Open Networking Foundation* (ONF). ONF is a consortium of software providers, content delivery networks, and networking equipment vendors whose purpose is to promote software-defined networking.

This discussion is based on the current OpenFlow specification, Version 1.3.0, June 25, 2012<sup>[8]</sup>. The original specification, 1.0, was developed at Stanford University and was widely implemented. OpenFlow 1.2 was the first release from ONF after inheriting the project from Stanford. OpenFlow 1.3 significantly expands the functions of the specification. Version 1.3 is likely to become the stable base upon which future commercial implementations for OpenFlow will be built. ONF intends for this version to be a stable target for chip and software vendors, so little if any change is planned for the foreseeable future<sup>[9]</sup>.

### Logical Switch Architecture

Figure 3 illustrates the basic structure of the OpenFlow environment. An SDN controller communicates with OpenFlow-compatible switches using the OpenFlow protocol running over the *Secure Sockets Layer* (SSL). Each switch connects to other OpenFlow switches and, possibly, to end-user devices that are the sources and destinations of packet flows. Within each switch, a series of tables—typically implemented in hardware or firmware—are used to manage the flows of packets through the switch.

Figure 3: OpenFlow Switch





The OpenFlow specification defines three types of tables in the logical switch architecture. A *Flow Table* matches incoming packets to a particular flow and specifies the functions that are to be performed on the packets. There may be multiple flow tables that operate in a pipeline fashion, as explained subsequently. A flow table may direct a flow to a *Group Table*, which may trigger a variety of actions that affect one or more flows. A *Meter Table* can trigger a variety of performance-related actions on a flow.

Before proceeding, it is helpful to define what the term *flow* means. Curiously, this term is not defined in the OpenFlow specification, nor is there an attempt to define it in virtually all of the literature on OpenFlow. In general terms, a flow is a sequence of packets traversing a network that share a set of header field values. For example, a flow could consist of all packets with the same source and destination IP addresses, or all packets with the same VLAN identifier. We provide a more specific definition subsequently.

### Flow-Table Components

The basic building block of the logical switch architecture is the flow table. Each packet that enters a switch passes through one or more flow tables. Each flow table contains entries consisting of six components:

- *Match Fields*: Used to select packets that match the values in the fields.
- *Priority*: Relative priority of table entries.
- *Counters*: Updated for matching packets. The OpenFlow specification defines a variety of timers. Examples include the number of received bytes and packets per port, per flow table, and per flow-table entry; number of dropped packets; and duration of a flow.
- *Instructions*: Actions to be taken if a match occurs.
- *Timeouts*: Maximum amount of idle time before a flow is expired by the switch.
- *Cookie*: Opaque data value chosen by the controller. May be used by the controller to filter flow statistics, flow modification, and flow deletion; not used when processing packets.

A flow table may include a *table-miss* flow entry, which renders all Match Fields wildcards (every field is a match regardless of value) and has the lowest priority (priority 0). The Match Fields component of a table entry consists of the following required fields:

- *Ingress Port*: The identifier of the port on the switch where the packet arrived. It may be a physical port or a switch-defined virtual port.
- *Ethernet Source and Destination Addresses*: Each entry can be an exact address, a bitmasked value for which only some of the address bits are checked, or a wildcard value (match any value).



- *IPv4 or IPv6 Protocol Number*: A protocol number value, indicating the next header in the packet.
- *IPv4 or IPv6 Source Address and Destination Address*: Each entry can be an exact address, a bitmasked value, a subnet mask value, or a wildcard value.
- *TCP Source and Destination Ports*: Exact match or wildcard value.
- *User Datagram Protocol (UDP) Source and Destination Ports*: Exact match or wildcard value.

The preceding match fields must be supported by any OpenFlow-compliant switch. The following fields may be optionally supported:

- *Physical Port*: Used to designate underlying physical port when packet is received on a logical port.
- *Metadata*: Additional information that can be passed from one table to another during the processing of a packet. Its use is discussed subsequently.
- *Ethernet Type*: Ethernet Type field.
- *VLAN ID and VLAN User Priority*: Fields in the IEEE 802.1Q Virtual LAN header.
- *IPv4 or IPv6 DS and ECN*: Differentiated Services and Explicit Congestion Notification fields.
- *Stream Control Transmission Protocol (SCTP) Source and Destination Ports*: Exact match or wildcard value.
- *Internet Control Message Protocol (ICMP) Type and Code Fields*: Exact match or wildcard value.
- *Address Resolution Protocol (ARP) Opcode*: Exact match in Ethernet Type field.
- *Source and Target IPv4 Addresses in Address Resolution Protocol (ARP) Payload*: Can be an exact address, a bitmasked value, a subnet mask value, or a wildcard value.
- *IPv6 Flow Label*: Exact match or wildcard.
- *ICMPv6 Type and Code fields*: Exact match or wildcard value.
- *IPv6 Neighbor Discovery Target Address*: In an IPv6 Neighbor Discovery message.
- *IPv6 Neighbor Discovery Source and Target Addresses*: Link-layer address options in an IPv6 Neighbor Discovery message.
- *Multiprotocol Label Switching (MPLS) Label Value, Traffic Class, and Bottom of Stack (BoS)*: Fields in the top label of an MPLS label stack.

Thus, OpenFlow can be used with network traffic involving a variety of protocols and network services. Note that at the MAC/link layer, only Ethernet is supported. Thus, OpenFlow as currently defined cannot control Layer 2 traffic over wireless networks.

We can now offer a more precise definition of the term *flow*. From the point of view of an individual switch, a flow is a sequence of packets that matches a specific entry in a flow table. The definition is packet-oriented, in the sense that it is a function of the values of header fields of the packets that constitute the flow, and not a function of the path they follow through the network. A combination of flow entries on multiple switches defines a flow that is bound to a specific path.

The *instructions component* of a table entry consists of a set of instructions that are executed if the packet matches the entry. Before describing the types of instructions, we need to define the terms “Action” and “Action Set.” Actions describe packet forwarding, packet modification, and group table processing operations. The OpenFlow specification includes the following actions:

- *Output*: Forward packet to specified port.
- *Set-Queue*: Sets the queue ID for a packet. When the packet is forwarded to a port using the output action, the queue id determines which queue attached to this port is used for scheduling and forwarding the packet. Forwarding behavior is dictated by the configuration of the queue and is used to provide basic QoS support.
- *Group*: Process packet through specified group.
- *Push-Tag/Pop-Tag*: Push or pop a tag field for a VLAN or MPLS packet.
- *Set-Field*: The various Set-Field actions are identified by their field type; they modify the values of respective header fields in the packet.
- *Change-TTL*: The various Change-TTL actions modify the values of the IPv4 Time To Live (TTL), IPv6 Hop Limit, or MPLS TTL in the packet.

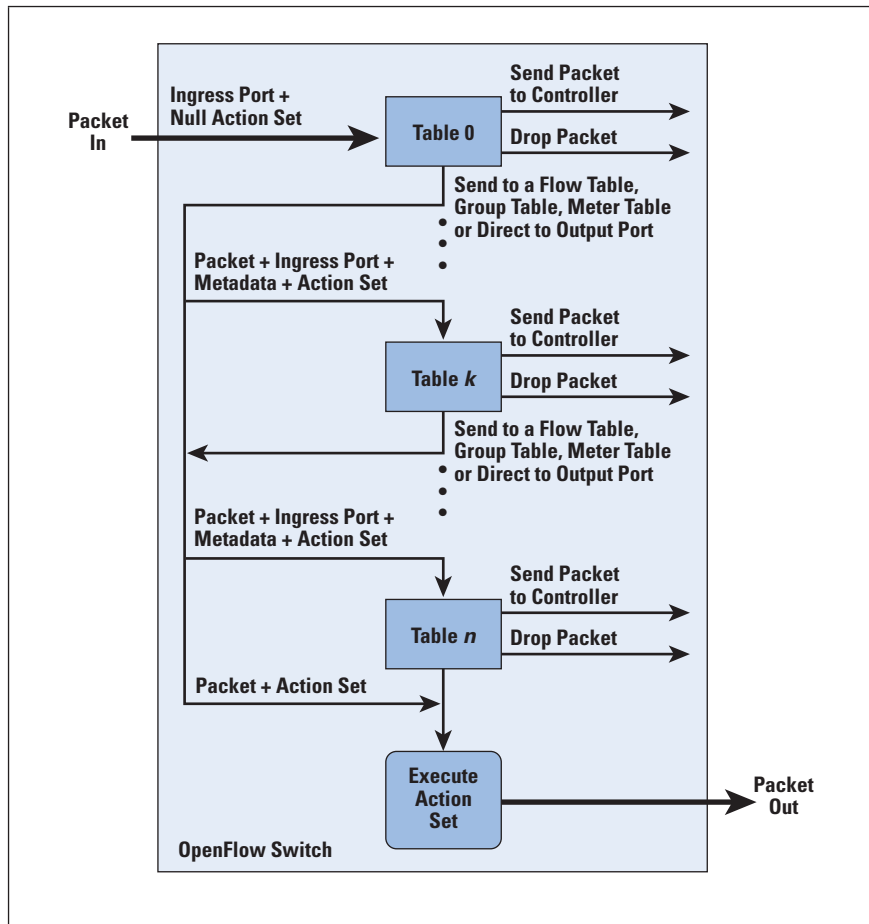
An *Action Set* is a list of actions associated with a packet that are accumulated while the packet is processed by each table and executed when the packet exits the processing pipeline. Instructions are of four types:

- *Direct packet through pipeline*: The Goto-Table instruction directs the packet to a table farther along in the pipeline. The Meter instruction directs the packet to a specified meter.
- *Perform action on packet*: Actions may be performed on the packet when it is matched to a table entry.
- *Update action set*: Merge specified actions into the current action set for this packet on this flow, or clear all the actions in the action set.
- *Update metadata*: A metadata value can be associated with a packet. It is used to carry information from one table to the next.

### Flow-Table Pipeline

A switch includes one or more flow tables. If there is more than one flow table, they are organized as a pipeline as shown in Figure 4, with the tables labeled with increasing numbers starting with 0.

Figure 4: Packet Flow Through OpenFlow-Compliant Switch



When a packet is presented to a table for matching, the input consists of the packet, the identity of the ingress port, the associated metadata value, and the associated action set. For Table 0, the metadata value is blank and the action set is null. Processing proceeds as follows:

1. Find the highest-priority matching flow entry. If there is no match on any entry and there is no table-miss entry, then the packet is dropped. If there is a match only on a table-miss entry, then that entry specifies one of three actions:
  - a. Send packet to controller. This action will enable the controller to define a new flow for this and similar packets, or decide to drop the packet.
  - b. Direct packet to another flow table farther down the pipeline.
  - c. Drop the packet.

2. If there is a match on one or more entries other than the table-miss entry, then the match is defined to be with the highest-priority matching entry. The following actions may then be performed:
  - a. Update any counters associated with this entry.
  - b. Execute any instructions associated with this entry. These instructions may include updating the action set, updating the metadata value, and performing actions.
  - c. The packet is then forwarded to a flow table further down the pipeline, to the group table, or to the meter table, or it could be directed to an output port.

For the final table in the pipeline, forwarding to another flow table is not an option.

If and when a packet is finally directed to an output port, the accumulated action set is executed and then the packet is queued for output.

### OpenFlow Protocol

The OpenFlow protocol describes message exchanges that take place between an OpenFlow controller and an OpenFlow switch. Typically, the protocol is implemented on top of SSL or *Transport Layer Security* (TLS), providing a secure OpenFlow channel.

The OpenFlow protocol enables the controller to perform add, update, and delete actions to the flow entries in the flow tables. It supports three types of messages, as shown in Table 1.

- *Controller-to-Switch*: These messages are initiated by the controller and, in some cases, require a response from the switch. This class of messages enables the controller to manage the logical state of the switch, including its configuration and details of flow- and group-table entries. Also included in this class is the Packet-out message. This message is used when a switch sends a packet to the controller and the controller decides not to drop the packet but to direct it to a switch output port.
- *Asynchronous*: These types of messages are sent without solicitation from the controller. This class includes various status messages to the controller. Also included is the Packet-in message, which may be used by the switch to send a packet to the controller when there is no flow-table match.
- *Symmetric*: These messages are sent without solicitation from either the controller or the switch. They are simple yet helpful. Hello messages are typically sent back and forth between the controller and switch when the connection is first established. Echo request and reply messages can be used by either the switch or controller to measure the latency or bandwidth of a controller-switch connection or just verify that the device is operating. The Experimenter message is used to stage features to be built into future versions of OpenFlow.

Table 1: OpenFlow Messages

Message		Description
Controller-to-Switch		
Features		Request the capabilities of a switch. Switch responds with a features reply that specifies its capabilities.
Configuration		Set and query configuration parameters. Switch responds with parameter settings.
Modify-State		Add, delete, and modify flow/group entries and set switch port properties.
Read-State		Collect information from switch, such as current configuration, statistics, and capabilities.
Packet-out		Direct packet to a specified port on the switch.
Barrier		Barrier request/reply messages are used by the controller to ensure message dependencies have been met or to receive notifications for completed operations.
Role-Request		Set or query role of the OpenFlow channel. Useful when switch connects to multiple controllers.
Asynchronous-Configuration		Set filter on asynchronous messages or query that filter. Useful when switch connects to multiple controllers.
Asynchronous		
Packet-in		Transfer packet to controller.
Flow-Removed		Inform the controller about the removal of a flow entry from a flow table.
Port-Status		Inform the controller of a change on a port.
Error		Notify controller of error or problem condition.
Symmetric		
Hello		Exchanged between the switch and controller upon connection startup.
Echo		Echo request/reply messages can be sent from either the switch or the controller, and they must return an echo reply.
Experimenter		For additional functions.

The OpenFlow protocol enables the controller to manage the logical structure of a switch, without regard to the details of how the switch implements the OpenFlow logical architecture.

### Summary

SDNs, implemented using OpenFlow, provide a powerful, vendor-independent approach to managing complex networks with dynamic demands. The software-defined network can continue to use many of the useful network technologies already in place, such as virtual LANs and an MPLS infrastructure. SDNs and OpenFlow are likely to become commonplace in large carrier networks, cloud infrastructures, and other networks that support the use of big data.

## References

- [1] Greg Goth, “Software-Defined Networking Could Shake Up More than Packets,” *IEEE Internet Computing*, July/August, 2011.
- [2] Robin Layland, “The Dark Side of Server Virtualization,” *Network World*, July 7, 2010.
- [3] Open Networking Foundation, “Software-Defined Networking: The New Norm for Networks,” ONF White Paper, April 12, 2012.
- [4] Dell, Inc., “Software Defined Networking: A Dell Point of View,” Dell White Paper, October 2012.
- [5] Hongtao Yin, Haiyong Xie, Tina Tsou, Diego Lopez, Pedro Aranda, and Ron Sidi, “SDNi: A Message Exchange Protocol for Software Defined Networks (SDNS) across Multiple Domains,” Internet Draft, work in progress, June 2012, **draft-yin-sdn-sdni-00.txt**
- [6] Steven Vaughan-Nichols, “OpenFlow: The Next Generation of the Network?” *Computer*, August 2011.
- [7] Thomas A. Limoncelli, “OpenFlow: A Radical New Idea in Networking,” *Communications of the ACM*, August 2012.
- [8] Open Networking Foundation, “OpenFlow Switch Specification Version 1.3.0,” June 25, 2012.
- [9] Sean Michael Kerner, “OpenFlow Protocol 1.3.0 Approved,” *Enterprise Networking Planet*, May 17, 2012.

WILLIAM STALLINGS is an independent consultant and author of many books on security, computer networking, and computer architecture. His latest book is *Data and Computer Communications* (Pearson, 2013). He maintains a computer science resource site for computer science students and professionals at **ComputerScienceStudent.com**. He has a Ph.D. in computer science from M.I.T. He can be reached at **ws@shore.net**

# IPv4 and IPv6 Address Authentication

by Scott Hogg, GTRI

Some Internet services use the source address of the client's computer as a form of authentication. These systems keep track of the *Internet Protocol* (IP) address that an end user used the last time that user accessed the site and try to determine if the user is legitimate. When that same user accesses the site from a different source IP address, the site asks for further authentication to revalidate the client's computer. The theory is that a user's typical location computer has a somewhat persistent IP address, but when the user has a new address, that user may be mobile or using a less secure wireless media, and then require further authentication. For example, many organizations have firewall policies with objects named like "Bob's Laptop" with the single IP address of his computer. This technique is used by some banking sites, some online gaming sites, and Gmail (for example, *Google Authenticator*)<sup>[1]</sup>.

Some online retailers track the client IP address for Business Intelligence or fraud detection and forensics purposes. The retailer tracks the client IP address using the source address to analyze fraudulent purchases and to track down criminal activity. Some industries frequently use the customer's IP address as a form of authentication. Also, many sites that use *Server Load Balancers* (SLBs) and *Application Delivery Controllers* (ADCs) use *X-forwarded-for* (XFF)<sup>[2]</sup> or *Hypertext Transfer Protocol* (HTTP) header insertion so that the back-end real servers are aware of the client's original IP address associated with the reverse proxy connection. The application can then use the IP address for tracking purposes or simply log the address with the transaction details.

Other applications try to validate the client's source IP address when the server receives an inbound connection. E-mail *Simple Mail Transfer Protocol* (SMTP) servers or *Internet Relay Chat* (IRC) servers can use the *Ident* protocol<sup>[3]</sup> to try to validate the originating e-mail server or client computer validity. SMTP e-mail servers<sup>[4]</sup> also use other protocols such as *SenderID*<sup>[5]</sup>, *Sender Policy Framework* (SPF)<sup>[6]</sup>, and *DomainKeys Identified Mail* (DKIM)<sup>[7]</sup> in an effort to restrict spam. *Domain Name System* (DNS) *pointer* (PTR) records are sometimes used as a way to confirm that the client IP address is configured in DNS (for example, forward-confirmed reverse DNS<sup>[8]</sup>).

Statically configured IP addresses are frequently used to signify some limited form of authentication. These addresses may not be used to authenticate a user, but authenticate IT systems to each other. Many manually configured systems rely on IP address to permit connectivity, including manually configured tunnels, *IP Security* (IPsec) peers, Apache *.htaccess*<sup>[9]</sup>, *.rhosts*<sup>[10]</sup>, SAMBA, and *Border Gateway Protocol* (BGP) peers, among many others.



The address is used as one part of the connection authentication. Obviously, IPsec connections are authenticated with certificates or preshared keys to strengthen their validation of the endpoints. Similarly, BGP peers use passwords (and/or *Time To Live* [TTL]<sup>[11]</sup>) to help secure the peer beyond just IP address confirmation.

Identity-based firewalls police users' network behavior by IP address through *Windows Active Directory*, *Remote Authentication Dial-In User Service* (RADIUS), or *Lightweight Directory Access Protocol* (LDAP). Palo Alto firewalls championed the *UserID* concept as part of their analysis of connections to permit or deny authentication<sup>[12]</sup>. The Cisco *Adaptive Security Appliance* (ASA) firewalls running Version 8.4 or later can be configured for Identify firewall functions<sup>[13]</sup>. Firewalls have always used manually configured IP addresses as the fundamental element of their policies. The IP address is used in the policy as if that concretely defines a system and/or user. This process of adding rules based on IP address continues until the firewall is a pincushion full of pinholes.

Organizations that rely on using an IP address as a form of authentication run the risk of an attacker learning that IP address and attacking using that address. Attackers who know the addresses that are being used could perform a *Man-in-the-Middle* (MITM) attack or use TCP session hijacking. The attacker needs to know only the information about which IP addresses are used for the communications. The attacker might be able to ascertain the IP addresses the organization uses by guessing or by other means. The attacker could find the external IP address of the company's firewall and assume that IPv4 *Network Address Translation* (NAT)<sup>[24]</sup> was being performed. The attacker could also suppose the business partner IP address. Organizations that use these techniques are relying on the secrecy of their IP addressing for the purposes of security.

### Address Quality

The quality of the IP address is an important concept to consider. For example, a global address is of higher surety and authenticity than a private address. Many organizations use private addresses and overlap between private networks, whereas global addresses are unique and they are registered to a specific entity. Public addresses can reveal the client's *Internet Service Provider* (ISP), the organization that has registered the IP addresses, and some geolocation information. However, any IP packet can be spoofed and the source-address modified or crafted. Of course, if the source IP addresses is spoofed, the return packets will not necessarily be sent back to the attacker's source in these cases, but one-way blind attacks are still possible. Furthermore, systems such as *Tor*<sup>[14]</sup> are intended to protect the identity of the end user.

Using the IP address as a form of authentication does not work if the client changes its location frequently. Today, many clients use mobile devices that can change their Layer 3 addresses often. The source IP address of the mobile device could change frequently and could even change during the transaction.

With increasing mobile device usage for business purposes, the ability to determine the typical IP address of the client becomes impossible. Increased scarcity of IPv4 addresses is leading service providers to use *Carrier-Grade NAT* (CGN) or *Large-Scale NAT* (LSN) and shorter and shorter *Dynamic Host Configuration Protocol* (DHCP) lease times, meaning that the client IP address is not static.

Many organizations and systems assume that a single computer with a single IP address represents a single user. The problem arises where IPv4 public addresses may not uniquely identify a single user. The industry may be trying to anticipate the implications of CGN/LSN and the effect of systems that rely on the uniqueness of a public IP address. Similar problems related to the mega-proxies of the late 1990s occurred (for example, AOL). With CGN/LSN systems in place, online retailers and banks will no longer be able to use the client IPv4 as the “real client IP.” Instead, the IP address observed on the retailer’s web servers will come from a pool of IPv4 addresses configured in the LSN system. In this situation, one bad actor could spoil that NAT pool IPv4 address for subsequent lawful users who follow. When a legitimate user tries to make an online purchase and that user’s system happens to use that IPv4 address of the bad actor, then the purchase attempt might be blocked. This situation would be bad for business on Cyber-Monday, or any day for that matter.

Table 1 compares IPv4 and IPv6 for their authentication purposes.

Table 1: IPv4 vs. IPv6 for Authentication

IPv4	IPv6
Extensive use of NAT	No motivation for NAT
End users use private addresses	End users use global addresses
Use of CGN/LSN starting	Abundance of IPv6 addresses
Robust geolocation	Geolocation needs improvement
Addresses could be spoofed	Addresses could be spoofed

### Public Addresses

Public IPv4 addresses are becoming increasingly scarce<sup>[15, 25]</sup>, however, an abundance of global IPv6 addresses are available<sup>[16]</sup>. Global IPv6 addresses can be obtained from *Regional Internet Registries* (RIRs) or from an IPv6-capable service provider. Residential broadband Internet users today use private IPv4 addresses on their internal computers, but these computers will soon start to use global IPv6 addresses as they upgrade to IPv6-capable *Customer Premises Equipment* (CPE). IPv6-enabled residential subscribers and employees of IPv6-enabled enterprises will be using global addresses when they access an IPv6-capable Internet service.

To online retailers, this situation may represent a change to their IP address authentication measures. As IPv4 residential users start to go through CGN/LSN systems, their IPv4 addresses will be useless for authentication.

However, their IPv6 addresses will be global addresses with no NAT taking place between the client and the server<sup>[17]</sup>. It will be seemingly more accurate to use the IPv6 address to determine the validity of the source. IPv6 could potentially help to create an environment with more “trustworthiness” and less anonymity. For example, IPv6 IPsec connections could use the *Authentication Header* (AH) and *Encapsulating Security Payload* (ESP) together to create stronger connections, where IPv4 IPsec connections rely on NAT-Traversal and can use only ESP<sup>[18]</sup>.

As we head toward an increasingly dual-stack world, applications will need to do “dual-checking” of both the client’s IPv4 and IPv6 addresses. In a dual-stack world, there is more work to do<sup>[19]</sup>, and servers using IP address authentication will need to understand that a single user will have both an IPv4 address and an IPv6 address and keep track of both. The other consideration is that IPv6 nodes may have multiple global IPv6 addresses in some situations.

### **Authentication with Addresses**

Security experts know that the secrecy of the encryption algorithm is not important, but the secrecy of the key is vitally important (Kerckhoffs’s Principle<sup>[20]</sup>). The same concept should hold true for an IP address. Users should not rely on the secrecy of their IP addresses to be secure; the security of the individual node should be strong enough to defend against attacks. To the extreme, users should feel confident enough in their security posture that they feel comfortable widely publicizing their IP address. However, even if you are using *LifeLock*<sup>[21]</sup>, you should still keep your Social Security Number or government ID number private.

Security practitioners know that authentication should involve multiple factors. A combination of “something you are” (biometrics), “something you know” (username/password) and “something you have” (token, *Common Access Card*<sup>[22]</sup>) forms a more solid foundation for identifying a user. Combining two factors provides more assurance than just one factor. We are all aware of the weaknesses of using username and password as a means of authentication<sup>[23]</sup>.

The systems mentioned so far in this article are three-factor systems (username, password, and IP address) which are presumably better than just username/password. However, we should acknowledge that an IP address is not a characteristic of a person. IP addresses have more to do with “somewhere you are,” because the IP address reflects location within a network topology by the prefix/subnet. The last few bits of an IPv4 address representing the point-of-attachment or an IPv6 *Interface Identifier* (IID) do not necessarily uniquely identify a user. Having authentication based on your location becomes difficult with mobile devices that roam widely. However, controlling authentication to users who are within the office subnet rather than outside the office may be useful.

An IP address is not something anyone really owns outright. Few organizations actually have complete ownership of their IP addresses. Organizations should read the fine print in the policies of their RIR. Organizations just pay RIR annual fees for their addresses, but if they stop paying those dues, the IP address allocation is revoked and the addresses go back into a pool for reallocation to another organization. Therefore, public IP addresses do not truly represent unequivocal ownership or legitimacy of a network.

### Conclusion

Many different types of systems use the client's source address as a form of authentication. Systems that rely on IP address checking will need to do so for IPv4 and will need to be modified to use IPv6 addresses. IPv6 systems will use global addresses without NAT, so the security systems must stand on their own even though the IPv6 address is publicized. IPv4 and IPv6 addresses can be spoofed, and as CGN/LSN systems become widely deployed the validity of a public IPv4 address decreases. However, IPv6 addresses are not necessarily any more trustworthy than IPv4 addresses when used for authentication. Regardless, the IP address should not be the only factor used for authentication, and we should not be using IPv4 or IPv6 addresses as a form of authentication. The truth is that the IT industry needs to be aware of where IP addresses are used as a form of authentication and seek out better forms of authentication beyond just username, password, and IP address.

### References

- [1] Google Authenticator,  
<http://support.google.com/a/bin/answer.py?hl=en&answer=1037451>
- [2] <http://en.wikipedia.org/wiki/X-Forwarded-For>
- [3] Mike St. Johns, "Identification Protocol," RFC 1413, February 1993.
- [4] [http://en.wikipedia.org/wiki/Email\\_authentication](http://en.wikipedia.org/wiki/Email_authentication)
- [5] Meng Weng Wong and Jim Lyon, "Sender ID: Authenticating E-Mail," RFC 4406, April 2006.
- [6] Wayne Schlitt and Meng Weng Wong, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1," RFC 4408, April 2006.
- [7] Miles Libbey, Michael Thomas, and Mark Delany, "DomainKeys Identified Mail (DKIM) Signatures," RFC 4871, May 2007.
- [8] [http://en.wikipedia.org/wiki/Forward-confirmed\\_reverse\\_DNS](http://en.wikipedia.org/wiki/Forward-confirmed_reverse_DNS)
- [9] <http://en.wikipedia.org/wiki/Htaccess>
- [10] <http://en.wikipedia.org/wiki/Rlogin>

- [11] Vijay Gill, John Heasley, and David Meyer, “The Generalized TTL Security Mechanism (GTSM),” RFC 3682, February 2004.
- [12] Palo Alto Networks, UserID,  
<http://www.paloaltonetworks.com/products/technologies/user-id.html>
- [13] Cisco ASA firmware 8.4 Identify Firewall,  
[http://www.cisco.com/en/US/docs/security/asa/asa84/configuration/guide/access\\_idfw.html](http://www.cisco.com/en/US/docs/security/asa/asa84/configuration/guide/access_idfw.html)
- [14] [http://en.wikipedia.org/wiki/Tor\\_\(anonymity\\_network\)](http://en.wikipedia.org/wiki/Tor_(anonymity_network))
- [15] [http://en.wikipedia.org/wiki/IPv4\\_address\\_depletion](http://en.wikipedia.org/wiki/IPv4_address_depletion)
- [16] <http://en.wikipedia.org/wiki/IPv6>
- [17] Scott Hogg and Owen DeLong, “IPv6 NAT - You can get it, but you may not need or want it,” Infoblox Blog, October 2, 2012.  
<http://www.infoblox.com/community/blog/ipv6-nat-you-can-get-it-you-may-not-need-or-want-it>
- [18] <http://en.wikipedia.org/wiki/Ipsec>
- [19] Scott Hogg, “Dual-Stack Will Increase Operating Expenses,” *Network World*, July 31, 2012,  
<http://www.networkworld.com/community/blog/dual-stack-will-increase-operating-expenses>
- [20] [http://en.wikipedia.org/wiki/Kerckhoffs%27s\\_Principle](http://en.wikipedia.org/wiki/Kerckhoffs%27s_Principle)
- [21] <http://en.wikipedia.org/wiki/LifeLock>
- [22] Common Access Card (CAC), <http://www.cac.mil/>
- [23] Mat Honan, “Kill the P@55W0rD,” *WIRED Magazine*, December 2012,  
<http://www.wired.com/gadgetlab/2012/11/ff-mat-honan-password-hacker/all/>
- [24] Geoff Huston, “Anatomy: A Look inside Network Address Translators,” *The Internet Protocol Journal*, Volume 7, No. 3, September 2004.
- [25] Several articles on IPv4 Exhaustion and IPv6 Transition in *The Internet Protocol Journal*, Volume 14, No. 1, March 2011.

SCOTT HOGG is the Director of Technology Solutions at GTRI in Denver Colorado. He holds a B.S. in Computer Science from Colorado State University, a M.S. in Telecommunications from the University of Colorado, and CCIE® #5133 and CISSP #4610 certifications. Hogg is active in the IPv6 community, Chair Emeritus of the RMv6TF, author of *IPv6 Security* (Cisco Press), a member of the Infoblox IPv6 Center of Excellence, a frequent presenter, and a *Network World* blogger. He can be reached at [scott@hoggnetwork.com](mailto:scott@hoggnetwork.com) or followed on twitter [@scotthogg](https://twitter.com/scotthogg)

# Summary Report of the ITU-T World Conference on International Telecommunications

by Robert Pepper and Chip Sharp, Cisco Systems

From 3–14 December, 2012, 151 Member States of the *International Telecommunication Union* (ITU) met in Dubai<sup>[0]</sup> at the *World Conference on International Telecommunications* (WCIT-12)<sup>[1]</sup> to revise the *International Telecommunication Regulations* (ITRs), a treaty-level document establishing policies governing international telecommunications services. During the 2-week conference the delegates debated several proposed changes on topics such as international mobile roaming, numbering, naming, addressing, fraud, the Internet, *Quality of Service* (QoS), etc. In the end, a revised version of the treaty was finalized<sup>[2]</sup>, but only 89 of the 151 Member States attending signed it.

There have been many articles discussing different aspects of the conference and its outcomes. This article provides background on the ITRs and focuses on the potential impact of the WCIT and its revised treaty on development of the Internet.

## Background

The ITRs originated from the development of international telegraphy in Europe in the late 1800s and the need for a treaty defining how the government-operated national telegraph networks would interconnect and interoperate<sup>[3]</sup>. As telephony and radio communications were invented, new treaties were developed to regulate their international operation. Up until the 1980s most telephone and telegraph companies were government-owned monopolies with some government licensed private companies operating as a monopoly. In 1988, the separate telegraph and telephone treaties were merged into the *International Telecommunications Regulations* while the *Radio Regulations* remained a separate treaty. By 1988, though some liberalization and privatization had started in a few countries in some regions, most international telecommunications services globally were still provided by monopoly, government-owned carriers, and services were dominated by voice rather than data. International Internet connectivity and traffic were practically nonexistent in most countries. Of course, international data traffic (including Internet) was growing in importance to some countries such as the United States and some large multinational companies such as IBM (which wanted to provide international *Virtual Private Networks* [VPNs]).

One important aspect of the ITRs in 1988 was the telephony accounting rate system. Briefly, this system consisted of a calling-party-pays business model for telephony in which the originating country pays the terminating country settlements based on a bilaterally agreed-upon accounting rate. Because developed countries tended to make more calls to developing countries than conversely and the accounting rate tended to be substantially above cost in many cases, the accounting rate system effectively became a subsidy program and a source for hard currency for developing countries.



Since 1988 market liberalization, reduced regulation, increased competition, and the rise of the Internet and mobile wireless industries have drastically changed the global communications landscape. In 1997, the U.S. *Federal Communications Commission* (FCC) opted out of the accounting rate system defined in the ITRs<sup>[4]</sup>, with many countries subsequently following suit. Voice over the Internet, arbitrage, hubbing, and other factors have reduced the telephony settlements revenue for developing countries. The 1988 ITRs<sup>[5]</sup> allowed for special arrangements between network operators outside the rules of the ITRs. These special arrangements allowed for the international physical connectivity on which growth of the international Internet depended.

As the Internet grew and the telecom market changed, there was increased pressure from some countries to revise the ITRs. Contributions submitted in the preparatory meetings for WCIT-12 reflected widely varying views on the nature and extent of possible changes to the ITRs to account for this greatly changed environment. Although some countries believed that the ITRs should set forth high-level strategic and policy principles that could adapt to further changes in the market, others proposed the inclusion of expanded regulatory provisions of a detailed and specific nature in the ITRs to address a wide range of new concerns and services, including the Internet, or even to include the intergovernmental regulation of content (for example, spam and information security).

#### High-Level Take-Aways

Out of 151 countries attending the conference, the treaty was signed by 89 countries, consisting of mostly emerging countries led by Russia, China, Brazil, and the Arab States; 55 countries, including the United States, Japan, Australia, Canada, United Kingdom, and most of Europe, did not sign at the time. Countries that did not sign the treaty in Dubai can accede to the treaty after the WCIT by notifying the Secretary-General of the ITU. It is quite likely that some countries that did not sign the treaty will accede to it over the next few years.

The treaty takes effect on January 1, 2015 (after the 2014 *Plenipotentiary Conference*). Each signing country has to go through its national process for approval (for example, ratification) before the treaty takes effect for that country.

Although there has been a lot of negative commentary on the WCIT in the Internet community, in the end there are some important positive results for the Internet:

- No provisions were added to treaty text explicitly concerning the Internet, Internet Governance, or information security.
- No provisions were added to the treaty text concerning naming or addressing.
- No provisions modifying the basic business models of the Internet or mandating QoS on the Internet were made.



- The updated treaty explicitly recognizes commercial arrangements in addition to the old accounting rate regime for telecommunications.
- Article 9 on *Special Arrangements* allowing for telecommunications arrangements outside the treaty was retained mostly unchanged, thus allowing such special arrangements to continue to be used even between nonsignatory and signatory countries.
- A new resolution on landlocked countries could encourage access of such countries to landing stations in other countries and ease landlocked countries' ability to acquire international connectivity.

Some results that could be of concern to the Internet follow:

- The term identifying the operators to which the treaty applies ("authorized operating agencies") was modified. The supporters of the new term claim it does not expand the scope of the treaty, but it will bear watching.
- A provision on "unsolicited bulk electronic communications," developed after a long debate on spam, could lead governments to regulate and filter e-mail in addition to having unintended consequences such as disallowing bulk electronic emergency warning systems.
- Numbering provisions and requirements to deliver *Calling Party Number* were intended by some countries to allow for restrictions on international *Voice over IP* (VoIP) and VoIP services (including VoIP over the Internet).
- A new provision on network security could encourage more multilateral discussions in an intergovernmental setting (as opposed to multistakeholder).
- A new Resolution 3 on the Internet instructs the Secretary-General to engage further in Internet Governance discussions and further supports intergovernmental Internet policy processes.
- A new Resolution 5 mentions the transition to IP-based networks. It originally was aimed at over-the-top providers, but was modified to apply to service providers of international services. The end result is rather ambiguous in many respects and will bear watching.
- A new Article was added concerning telecommunication exchange points. Although the Internet is not mentioned explicitly, the originators of this article intended for it to apply to Internet Exchange Points. This Article could be used to support development of an enabling environment for regional telecommunication connectivity, but could also be used to justify regulation of Internet Exchange Points.
- Resolution Plen/4 requires PP'14 to consider a review of the ITRs every 8 years. This provision could result in another WCIT in 2020.

Table 1 lists the Member States that signed and did not sign the treaty in Dubai<sup>[6]</sup>.

*Table 1: Treaty Signatories and Nonsignatories*

Signatories			Nonsignatories	
Afghanistan	Guatemala	Qatar	Albania	Latvia
Algeria	Guyana	Russia	Andorra	Lichtenstein
Angola	Haiti	Rwanda	Armenia	Lithuania
Argentina	Indonesia	Saint Lucia	Australia	Luxembourg
Azerbaijan	Iran	Saudi Arabia	Austria	Malawi
Bahrain	Iraq	Senegal	Belarus	Malta
Bangladesh	Jamaica	Sierra Leone	Belgium	Marshall Islands
Barbados	Jordan	Singapore	Bulgaria	Moldova
Belize	Kazakhstan	Somalia	Canada	Mongolia
Benin	Korea (Rep. of)	South Africa	Chile	Montenegro
Bhutan	Kuwait	South Sudan	Colombia	Netherlands
Botswana	Kyrgyzstan	Sri Lanka	Costa Rica	New Zealand
Brazil	Lebanon	Sudan	Croatia	Norway
Brunei	Lesotho	Swaziland	Cyprus	Philippines
Burkina Faso	Liberia	Tanzania	Czech Republic	Poland
Burundi	Libya	Thailand	Denmark	Peru
Cambodia	Malaysia	Togo	Estonia	Portugal
Cape Verde	Mali	Trinidad and Tobago	Finland	Serbia
Central African Rep.	Mauritius	Tunisia	France	Slovak Republic
China	Mexico	Turkey	Gambia	Slovenia
Comoros	Morocco	Uganda	Georgia	Spain
Congo	Mozambique	Ukraine	Germany	Sweden
Cote d'Ivoire	Namibia	UAE	Greece	Switzerland
Cuba	Nepal	Uruguay	Hungary	United Kingdom
Djibouti	Niger	Uzbekistan	India	United States
Dominican Rep.	Nigeria	Venezuela	Ireland	
Egypt	Oman	Vietnam	Israel	
El Salvador	Panama	Yemen	Italy	
Gabon	Papua New Guinea	Zimbabwe	Japan	
Ghana	Paraguay		Kenya	

**Note:** Other *United Nations* (UN) member states were not eligible to sign or did not attend the conference but might still accede to the treaty: Antigua and Barbuda, Bahamas, Bolivia, Bosnia and Herzegovina, Cameroon, Chad, Dem. People’s Republic of Korea, Dem. Rep. of the Congo, Dominica, Ecuador, Equatorial Guinea, Eritrea, Ethiopia, Fiji, Grenada, Guinea, Guinea-Bissau, Honduras, Iceland, Kiribati, Lao P.D.R., T.F.Y.R. Macedonia, Madagascar, Maldives, Mauritania, Micronesia, Monaco, Myanmar, Nauru, Nicaragua, Pakistan, Romania, Saint Kitts and Nevis, Saint Vincent and the Grenadines, Samoa, San Marino, Sao Tome and Principe, Seychelles, Solomon Islands, Suriname, Syria, Tajikistan, Timor-Leste, Tonga, Turkmenistan, Tuvalu, Vanuatu, the Vatican, and Zambia.

### Proposals and Outcomes

When the conference began there were several provisions that either explicitly or implicitly applied to the Internet, including:

- A proposal to define the term “Internet” and explicitly bring the Internet into the regulatory structure of the treaty
- Proposals to bring Internet naming, addressing, and identifiers into the treaty
- A proposal to include a provision on access to Internet websites
- A proposal on “traffic exchange points” that was intended to apply to *Internet Exchange Points*
- Proposals from multiple states on spam, information security, and *cybersecurity*

Although the Secretary-General of the ITU declared that the WCIT was not about the Internet or Internet Governance<sup>[7]</sup>, by rule, the WCIT had to consider input from its Member States. Given that Member States submitted proposals on the Internet, the Internet and Internet Governance was a substantive topic of discussion.

The following sections provide a brief review of some of the more difficult discussions related to the Internet.

### Security

There were several proposals<sup>[8]</sup> going into the WCIT to include cybersecurity, including information security, in the new ITRs. These proposals generated significant discussions and negotiations during the conference. The final text (Article 5A) is a great improvement over the proposals into the conference in that it focuses on the security and robustness of networks and prevention of technical harm to networks, with no mention of information security or cybersecurity.

The new provision mentions that Member States shall “collectively endeavour,” a provision that could engender more multilateral discussions in an intergovernmental setting (for example, ITU).

### Organizations to Which the Treaty Applies

The 1988 ITR treaty focused on licensed carriers and government-owned *Post, Telephone, and Telegraph* (PTT) entities. Proposals<sup>[8]</sup> into the WCIT would have applied the treaty to a wider range of organizations and companies. In the end, the treaty developed a new term, *Authorized Operating Agencies* (AOA). The proponents of this new term argued that it does not broaden the scope of the ITRs in terms of the organizations to which it applies. This interpretation of the new term should be supported, but monitored.

### Internet-Specific Proposals and Resolutions (Resolutions Plen/3 and Plen/5)

Proposals<sup>[8]</sup> were submitted to the WCIT to define the term “Internet” and to encode into the treaty the right of countries to regulate the “national segment” of the Internet. At the end of the first week of WCIT, Algeria, Saudi Arabia, Bahrain, China, United Arab Emirates, Iraq, Sudan, and Russia announced development of a new draft set of Resolutions that contained provisions that Member States shall have the right to manage the Internet, including Internet numbering, naming, addressing, and identification resources.

Although the United States, United Kingdom, and others were successful in removing any mention of the Internet from the treaty text, Internet-related language was moved into a nonbinding resolution (*Resolution Plen/3*) proposed by Russia “to foster an enabling environment for the greater growth of the Internet.” Resolution Plen/3 instructs the ITU Secretary-General “to continue to take the necessary steps for ITU to play an active and constructive role in the development of broadband and the multistakeholder model of the Internet as expressed in § 35 of the *Tunis Agenda*.” It also invites Member States to elaborate their positions on Internet-related concerns in the relevant ITU-related fora (something they could have done anyway).

This does not look too bad until one reads Paragraph 35 of the Tunis Agenda<sup>[9]</sup>. This paragraph lays out the roles of each type of stakeholder (private industry, civil society, *Intergovernmental Organizations* [IGOs], governments, etc.). It reserves an explicit role in “Internet-related public policy issues” for governments and intergovernmental organizations. It does not provide for any role in this area for the private sector or civil society. So although the Resolution seems to support the multistakeholder model of the Internet, it really restricts the roles of several of the main stakeholders.

Several countries pushed for inclusion of Paragraph 55 of the Tunis Agenda, recognizing that the existing arrangements have worked effectively, to balance the inclusion of Paragraph 35, but it was not included in the final Resolution.

Resolution Plen/3 may be used by some governments to reinforce the ITU’s role in Internet Governance, including at future ITU conferences in 2013 and 2014.

On the other hand, the Resolution also instructs the Secretary-General “to support the participation of Member States and *all* other stakeholders, as applicable, in the activities of ITU in this regard.” This statement supports participation of all stakeholders in the activities of the ITU, not restricted just to ITU Members, or in the case of ITU Council or some Council Working Groups just to Member States.

In signing the Final Acts, Russia added a Declaration/Reservation that it views the Internet as a new global telecommunication infrastructure and reserves the right to implement public policy, including international policy, on matters of Internet Governance. This reservation could signal that Russia plans to apply the telecommunications provisions in the ITRs to the Internet and to further regulate the Internet.

In addition to Resolution Plen/3, some of the proposals on the Internet were part of the discussion on Resolution Plen/5. This Resolution began as a basic resolution on invoicing for international telecommunication services, but ended up including numerous other provisions that did not make it into the main text of the treaty. Although the final text does not contain provisions explicitly mentioning the Internet, the introductory text of the Resolution mentions the transition of phone and data networks to IP-based networks. Also, the proposal that evolved into “resolves” originally applied to the relationship between network operators and application providers. During the discussions this proposal was modified to refer to “providers of international services” instead of application providers. Even with this modification, the application of this provision is ambiguous and could be applied to over-the-top providers.

Resolution Plen/5 is likely to reinforce work in Study Group 3 on accounting, fraud and charges for international telecommunications service traffic termination and exchange, etc.

#### **Telecommunications Traffic Exchange Points**

A proposal<sup>[8]</sup> concerning “telecommunication traffic exchange points” was included as an Article in the ITRs. The term “telecommunication traffic exchange point” was left undefined. This article does not mention the Internet or Internet Exchange Points, but the discussion of this Article included discussion on how it related to Internet Exchange Points. At least one delegation indicated that the Article was intended to help enable development of regional Internet Exchange Points.

Although this provision raised concerns over possible regulation of Internet Exchange Points, it focuses on creating an enabling environment for creation of regional telecommunication traffic exchange points. This environment could provide support for development of trans-border telecommunications and connectivity.

### **Route-Related Factors**

Prior to the conference, there were several proposals [8] to require transparency into the international routes used for a Member States' traffic and to allow Member States to control what routes were used between them. Note that the definition of "route" in the ITRs is different from the concept of a "route" on the Internet. In the ITRs, a route is defined as the technical facilities used for telecommunications traffic between two telecommunication terminal exchanges or offices.

Coming into the WCIT, the proposal to control routes by Member States was dropped from the proposal, so the debate centered over whether Member States should have the right to know what routes were being used. After much discussion, the final result was a provision allowing "authorized operating agencies" (not Member States) to determine which routes are to be used between them and allowing the originating operator to determine the outbound route for traffic. This provision is not much different from how network operators manage their networks today.

### **Quality of Service Proposals**

Several proposals<sup>[8]</sup> were made to WCIT to require QoS to be negotiated between network providers including Internet providers. Some proposals also allowed network providers to charge over-the-top providers for QoS.

The final provisions did not add any new requirements for QoS other than a nonspecific requirement related to mobile roaming. Although no new provisions were added specific to the Internet, it does not mean that countries could not try to impose the current QoS provisions to VoIP services. The debate over QoS on the Internet will continue outside the ITRs.

### **Naming, Numbering, and Addressing Proposals**

Several countries and regions proposed<sup>[8]</sup> to extend provisions on telephone numbering to include naming, addressing, and origin identifiers. Several proposals were made to require delivery of calling party number and to cooperate in preventing the misuse ("misuse" not defined) of numbering, naming, and addressing resources. Although the Internet was not explicitly mentioned, these proposals were intended to apply to VoIP based on comments at pre-WCIT preparatory meetings<sup>[10]</sup>.

In the end, several provisions were added related to delivery of calling party number and prevention of misuse of telecommunications numbering resources as defined in ITU-T Recommendations. Provisions to include naming, addressing, and more general "origin identifiers" were not accepted.

Even though there were no provisions specifically on the Internet, some countries could apply these provisions to VoIP services that use E.164 telephone numbers and that provide for bypass of the international telephony accounting system.

However, it is not clear that these provisions add any more authority than what these countries have today.

### **Content and Spam**

Proposals<sup>[8]</sup> to include spam in the treaty caused a lot of contentious discussion, in ad hoc groups, plenary, and in consultations. Some countries took a strong position that spam is a content topic that was out of scope of the ITRs. There was a concern that adding a provision on spam would legitimize content filtering by governments. Some African countries insisted on including a provision on spam, claiming that it consumed a large percentage of their international bandwidth. In the end to address concern about content, a statement was added to Article 1.1:

“These Regulations do not address the content-related aspects of telecommunications.”

To address the proposals on spam, Article 5B was added on unsolicited bulk electronic communication:

“Member States should endeavour to take necessary measures to prevent the propagation of unsolicited bulk electronic communications and minimize its impact on international telecommunication services. Member States are encouraged to cooperate in that sense.”

As written the final text, it is fairly vague and could have implications beyond spam; for example, there are no exemptions for broadcasters or for emergency alert systems (for example, tsunami alerts). It is also not clear how Article 5B can be implemented consistent with the statement on content in Article 1.1.

It was clear from the discussion that many of the delegates from countries supporting this provision do not understand spam or spam-mitigation techniques and their usage (or not) in their own countries. It is clear that many of the delegates were not aware of basic best practices from the *Messaging Anti-Abuse Working Group* (MAAWG) and other organizations. These discussions highlighted the need for capacity building for developing countries on spam-mitigation techniques.

### **Human Rights and Member State Access to International Telecommunications**

In a plenary session on the penultimate night of the WCIT, a provision on human rights was added to the final draft of the ITRs. This discussion led to a debate concerning the right of Member States to access international telecommunication services, originating from a proposal from Sudan and Cuba creating a right of Member States to access Internet websites. This provision was targeted at U.S. and European actions taken in response to UN sanctions against Sudan due to Darfur and U.S. sanctions on Cuba.



The provision provides a right for Member States, not its citizens. Thus it did not provide any rights for citizens to access international telecommunication services. In addition, it is not clear what or whose international telecommunication service Member States have a right to. The implications of the provision were unclear, and delegations did not have time to consult their home countries before the end of the conference.

Several times during the debate the Chair of the WCIT and the Secretary-General of the ITU both tried to dissuade the proponents from pushing their proposal, to no avail. After extended debate, Iran called for a point of order and then called for a vote, the only official vote of the conference. After the text passed by majority vote, the Chair of the WCIT declared the ITRs approved. At that point the United States, followed by the United Kingdom, Sweden, and other countries, made statements that they would not sign the treaty. Supporters of the treaty read their statements in favor of the treaty. The conference was effectively over<sup>[11]</sup>.

The uncertainty caused by the addition of this text at such a late date and the way it was added created a situation in which many countries that might have signed the treaty ended up not signing. This provision more than any other disagreement in the conference caused the conference to split to the extent that it did.

### Looking Forward

Much of the long-term impact of the treaty will not be felt until the signing governments ratify the treaty and start enacting provisions into either law or regulation. It is likely that some of the countries that did not sign in Dubai will accede to the treaty at a later time, including countries that did not attend the WCIT.

WCIT is only one step (though an important one) in the long-term debate over Internet Governance and the appropriate role of governments (and intergovernmental organizations) in the Internet. The debate will continue in numerous international fora going forward such as:

- World Telecommunications Policy Forum (May 2013)
- World Summit on the Information Society Action Line Forum (May 2013)
- ITU Council Working Group on Internet Public Policy (ongoing)
- ITU-T Study Group meetings (ongoing)
- ITU Plenipotentiary Conference (2014)
- WSIS+10 Review (2013–2015)

It has already been seen that many of the same topics debated at WCIT will be debated in these venues; for example, IP addressing, naming, spam, and cybersecurity. The WCIT Resolutions (especially Res. Plen/3) will likely be used to promote a larger role of the ITU in the Internet Governance debate.

The ITU's Plenipotentiary Conference in 2014 will be the next important treaty conference where the ITU's Constitution and Convention (both treaty instruments) can be revised. In the hierarchy of treaties at ITU, the ITU Constitution takes precedence over the ITRs, and many of the terms used in the ITRs are defined in the Constitution. Therefore, changes to the ITU Constitution could affect the meaning of the ITRs. The ITU Plenipotentiary will provide an opportunity for the ITU Member States to come together and heal some of the differences coming out of the WCIT, but it is also an opportunity to widen the rift.

The WSIS+10 Review will be an important process because it is likely to set the agenda for the discussion of Internet Governance for the 5–10 years after 2015, much as the Tunis Agenda from 2005 set the agenda for the last 8 years. An important aspect of the WSIS+10 Review is that it involves other UN agencies (for example, UNESCO) in addition to the ITU. Many of the events involve stakeholders whose voices are not normally heard at ITU conferences.

Some of the disagreements exhibited at WCIT brought to light opportunities for the Internet community to engage with governments and other stakeholders by providing technical and thought leadership. Capacity building with many of the developing country governments will be an important part of the preparation leading up to the major international conferences such as the ITU Plenipotentiary and WSIS+10.

Much of the growth of the Internet going forward is likely to come in the countries that signed the ITRs. Many of these countries have started to develop multistakeholder consultations and processes when dealing with Internet topics. The fact that a government signed the ITRs does not mean that the country is somehow against the Internet. On the contrary, many of these countries are looking for ways to accelerate the Internet's development within their borders and to accelerate their international connectivity to the Internet. As the Internet grows and develops in these countries, the Internet communities in these countries will likely look to play a larger role in a consultative process regarding government positions on issues related to Internet Governance. Future growth of the Internet across ITR boundaries (signatories and non-signatories) will depend on cooperation amongst all stakeholders.

## References

- [0] Geoff Huston, “December in Dubai: Number Misuse, WCIT, and ITRs,” *The Internet Protocol Journal*, Volume 15, No. 2, June 2012.
- [1] World Conference on International Telecommunications (WCIT-12),  
<http://www.itu.int/en/wcit-12/Pages/default.aspx>
- [2] International Telecommunication Regulations,  
<http://www.itu.int/en/wcit-12/Pages/itrs.aspx>
- [3] “Discover ITU’s History,” <http://www.itu.int/en/history/Pages/DiscoverITUsHistory.aspx>
- [4] “International Settlements Policy and U.S.–International Accounting Rates,”  
<http://www.fcc.gov/encyclopedia/international-settlements-policy-and-us-international-accounting-rates>
- [5] “WATTC-88 World Administrative Telegraph and Telephone Conference (Melbourne, 1988),”  
<http://www.itu.int/en/history/Pages/TelegraphAndTelephoneConferences.aspx?conf=33&dms=S0201000021>
- [6] “Signatories of the Final Acts: 89 (in green),”  
<http://www.itu.int/osg/wcit-12/highlights/signatories.html>
- [7] “Dr. Hamadoun I. Touré, ITU Secretary-General First Plenary of World Conference on International Telecommunications (WCIT-12),”  
<http://www.itu.int/en/wcit-12/Pages/speech-toure2.aspx>
- [8] “Proposals Received from ITU Member States for the Work of the Conference,”  
[http://www.itu.int/md/dologin\\_md.asp?lang=en&id=S12-WCIT12-121203-TD-0001!!MSW-E](http://www.itu.int/md/dologin_md.asp?lang=en&id=S12-WCIT12-121203-TD-0001!!MSW-E)
- [9] “Tunis Agenda for the Information Society,” *World Summit on the Information Society*, 2005. <http://www.itu.int/wsis/docs2/tunis/off/6rev1.html>
- [10] Council Working Group to Prepare for the 2012 WCIT,  
<http://www.itu.int/council/groups/cwg-wcit12/index.html>
- [11] Webcast and Captioning of the WCIT,  
<http://www.itu.int/en/wcit-12/Pages/webcast.aspx>

ROBERT PEPPER leads Cisco's Global Technology Policy team working with governments across the world in areas such as broadband, IP enabled services, wireless and spectrum policy, security, privacy, Internet governance and ICT development. He joined Cisco in July 2005 from the FCC where he served as Chief of the Office of Plans and Policy and Chief of Policy Development beginning in 1989 where he led teams developing policies promoting the development of the Internet, implementing telecommunications legislation, planning for the transition to digital television, and designing and implementing the first U.S. spectrum auctions. He serves on the board of the U.S. Telecommunications Training Institute (USTTI) and advisory boards for Columbia University and Michigan State University, and is a Communications Program Fellow at the Aspen Institute. He is a member of the U.S. Department of Commerce's Spectrum Management Advisory Committee, the UK's Ofcom Spectrum Advisory Board and the U.S. Department of State's Advisory Committee on International Communications and Information Policy. Pepper received his BA. and Ph.D. from the University of Wisconsin-Madison. E-mail: [\*\*rmpepper@cisco.com\*\*](mailto:rmpepper@cisco.com)

CHIP SHARP has 30 years in the communications industry and currently is a Director in the Research and Advanced Development Department at Cisco Systems, Inc. His current role is in Technology Policy focusing on Internet Governance issues. He participated in the US preparatory process for WCIT from 2010 including as Private Sector Advisor to the US Delegation to the ITU's Council Working Group to Prepare for the 2012 WCIT (CWG-WCIT12), mainly analyzing the impact of proposals on the Internet and Internet Governance. He helped develop many of the talking points and position papers related to the Internet for the US Delegation. He served in the same capacity on the US Delegation to WCIT. He continues to be active in many follow-on activities preparing for the World Telecommunication Policy Forum (WTPF), World Summit on the Information Society 10 year review (WSIS+10), ITU Plenipotentiary Conference 2014 and Internet Governance Forum. He also currently service on the FCC's Open Internet Advisory Committee. Prior to this role, he led a multinational, multidisciplinary team at Cisco helping drive various technologies such as LISP, DNSSEC, BGPSEC, ENUM, Lawful Intercept etc. He has also supported capacity building and development programs for developing countries, for example, deployment of Internet Exchange Points (IXPs). He started at Cisco in 1996 helping design dialup Internet access products to interface with legacy telco signaling systems. Prior to Cisco, Chip worked at Teleos Communications, AT&T Consumer Product Labs and NASA's Communications Division. E-mail: [\*\*chsharp@cisco.com\*\*](mailto:chsharp@cisco.com)

## Letters to the Editor

Dear Ole,

I am sorry that there is some delay (more than 1 second) between the arrival of *The Internet Protocol Journal* at my desk and this e-mail. In the December 2012 issue (Volume 15, No. 4), Geoff Houston discusses the extra second on the last minute of the 31st of June. There is no 31st of June in the calendar, at least not in old Europe, but maybe in the United States. It is funny to discuss the problem of a second at the end of a nonexistent day, isn't it?

Nevertheless I could take some new knowledge from this article.

Best regards,

—Richard Schuerger  
`richard.schuerger@gmx.de`

Hi Geoff (and Ole)!

I am sitting comfortably in a chair on the terrace in a Tenerife house, reading the December 2012 issue of IPJ, which I received by mail today. Since I have been working many years with the *Network Time Protocol* (NTP), I started reading your article on the subject with great interest. Having read only a few sentences I jumped in my chair:

“Back at the end of June 2012 there was a brief IT hiccup as the world adjusted the *Coordinated Universal Time* (UTC) standard by adding an extra second to the last minute of the 31st [!!] of June.”

Of course you may have received numerous notices of this hiccup [ha, ha], but still I couldn't resist writing to you. Thank you for an [otherwise] well-written and clarifying article (as always).

—Truls Hjelle  
`truls@sund-hjelle.org`

PS: Thanks to Ole for this anachronism on paper still available to us oldies who prefer sitting with a paper magazine in the sun instead of gazing at a poorly lit screen and struggling with the tiny letters.


*The author responds:*

Back in 45 BC, Julius Caesar made same revolutionary changes to the Roman calendar, and the changes included adding one extra day to June (well not quite, as the letter “J” was not around until the 16th Century, and the letter “u” was also yet to make its debut, so it is probably less of an anachronism to record that Gaius Iulius Caesar added an extra day to the month of Iunius). Either way, this change brought the total number of days in the month of June to 30, which is where it has remained for 2058 years.

It is often said that Australia operates on a calendar all of its own, but while our isolation on a largish rock at the southern end of the Pacific Ocean has led to a number of revolutionary innovations that are easily on a par with fire and the wheel, including the world-renowned stump-jump plough and the sheep-shearing machine, we Australians have not yet turned our collective national genius to the calendar. Despite a pretty sensible suggestion from the latest meeting of the Grong Grong Shire Council for a year to be made up of 10 months of 30 days followed by a decent 65-day session at the pub, we have yet to get the blokes back from the pub after their last 65-day bender, so that plan needs some more work back at the shed before it gets another airing! Thus it looks like Australia uses the same calendar as everyone else, making the reference to the 31st of June one of those pesky brain-fade errors! Oops. Yes, it was meant to say 30th of June. Well spotted!

—*Geoff Huston*

**gih@apnic.net**

 The Internet Protocol Journal, Cisco Systems 170 West Tasman Drive San Jose, CA 95134-1706 USA ADDRESS SERVICE REQUESTED	<div>PSRST STD U.S. Postage PAID PERMIT No. 5187 SAN JOSE, CA</div> <div>      FOSTER BUNNY LAGOMORPH INC 1234 MAIN STREET SAN FRANCISCO, CA 94104-1234</div>
SUBSCRIPTION ID: FBUNN006188      V15 N3 EXPIRATION DATE: 15-JAN-2013	

Don't forget to renew and update your subscription. For details see the IPJ Subscription FAQ in our previous issue (Volume 15, No. 4).

## Book Review

### On Internet Freedom

*On Internet Freedom*, by Marvin Ammori, Elkat Books, January 2013, sold by: Amazon Digital Services, Inc., ASIN: B00B1MQZNW.

Marvin Ammori has written an important book about the threats to free speech and expression that we are not only privileged to conduct on the Internet today but have come to treat as basic human rights.

*On Internet Freedom* looks at the past, present, and future of the Internet as a speech technology. Ammori examines how the coordinated and determined efforts by Big Content to protect content and increasing efforts by governments to censor content threaten Internet use as we embrace it today. Ammori also explains how these acts were in fact anticipated by Clark, Sollins, Wroclawski, and Braden in a paper entitled “Tussle in Cyberspace: Defining Tomorrow’s Internet,”<sup>[1]</sup> where the authors assert:

“User empowerment, to many, is a basic Internet principle, but for this paper, it is the manifestation of the right to choose—to drive competition, and thus drive change.”

Ammori cites only the first clause of this sentence—as a technologist, I believe the second is extremely important as well—but he makes clear that the end-to-end design of the Internet establishes a fundamental thesis:

“If user choice is our design principle, then users should have the final say.”

Unfortunately, Ammori explains that users do not have the final say but are increasingly challenged by lawyers, bureaucrats, commissioners, and others who are motivated to constrain their freedoms and who want to do so by altering the fundamental design of the Internet. Ammori’s response, admittedly U.S.-centric, is simple: the Internet is a speech technology, and:

“... the ultimate design principle for any speech technology, at least in the United States: the First Amendment, which protects freedom of speech. The *First Amendment* is not generally thought of as a design principle, but, by definition, it limits what Congress or any other government actor may or may not adopt in shaping the Internet’s future.”

This statement sets the context for the remainder of the book. In Part II, Ammori looks at events leading to the 18 January 2012 Internet Blackout in protest of the *Stop Online Piracy Act* (SOPA) and *PROTECT IP Act* (PIPA) and how these and possibly future legislation threaten “...the speech tools of the many while reshaping our speech environment for the benefit of the few.”



Conveniently, Part II is largely about how the few benefit. Before judging whether you believe this theory is even-handed or not, remember that the litmus test throughout this book is the First Amendment of the U.S. Constitution. This part ought to make every Internet user or free speech advocate pause, or shiver. One of the most worrisome speculations Ammori offers is the extent to which legislation could stilt adoption of emerging technologies such as *three-dimensional* (3D) printing or stifle future innovations of this kind.

Part III looks at how the Internet as speech technology influences governments, how governments have attempted to exert influence, and how Internet users and dominant Internet forces (Google, Amazon, Facebook, and Twitter) respond. This part will probably be illuminating for most readers, because it explains situations where a *private conversation* between a government official and an *Internet Service Provider* (ISP) or hosting company can circumvent the First Amendment, and why *Terms of Service* are often more speech-restricting than the First Amendment as well.

Part IV focuses on net neutrality concerns. Ammori draws the lines of conflict: ISPs seek to differentiate, rate-control, block, or charge users differently for content that is transmitted on their networks. However, content includes speech, and if the Internet is speech technology, then ISPs should not be able to decide what you say or see, or they do so in violation of your First Amendment rights. Ammori also explains that net neutrality is not only a First Amendment concern but also an economic one: net neutrality violations can influence investments in or creation of new technology.

I began by saying that Marvin Ammori has written an important book. It is also an extremely readable book. Ammori does a commendable job explaining constitutional law and technology in easy to understand terms. I highly recommend the book not only for people who are interested in law or technology but for anyone who advocates freedom of expression.

On Internet Freedom is currently available as a Kindle download.

- [1] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," *IEEE/ACM Transactions on Networking*, Volume 13, Issue 3, June 2005. Available from:  
<http://groups.csail.mit.edu/ana/Publications/PubPDFs/Tussle2002.pdf>

—Dave Piscitello, [dave@corecom.com](mailto:dave@corecom.com)

Reprinted with permission from *The Security Skeptic* blog:  
<http://securityskeptic.typepad.com/the-security-skeptic/>

### Nominations Sought for 2013 Jonathan B. Postel Service Award

The Internet Society is soliciting nominations of qualified candidates for the 2013 *Jonathan B. Postel Service Award* by May 31, 2013. This annual award is presented to an individual or organization that has made outstanding contributions in service to the data communications community. The award is scheduled to be presented during the 87th IETF meeting in Berlin, Germany, July 28–August 2.

The award was established by the Internet Society to honor a person who has made outstanding contributions in service to the data communications community. The award is focused on sustained and substantial technical contributions, service to the community, and leadership. With respect to leadership, the award committee places particular emphasis on candidates who have supported and enabled others in addition to their own specific actions.

The award is named for Dr. Jonathan B. Postel to recognize and commemorate the extraordinary stewardship exercised by Jon over the course of a thirty-year career in networking. He served as the editor of the RFC series of notes from its inception in 1969 until 1998. He also served as the ARPANET “Numbers Czar” and *Internet Assigned Numbers Authority* (IANA) over the same period of time. He was a founding member of the Internet Architecture (nee Activities) Board and the first individual member of the Internet Society, which he also served as a Trustee.

For more information, see: <http://www.internetsociety.org/>

### Upcoming Events

The *Internet Corporation for Assigned Names and Numbers* (ICANN) will meet in Beijing, China, April 7–11, 2013 and in Durban, South Africa, July 14–18, 2013. For more information, see: <http://icann.org/>

The *North American Network Operators’ Group* (NANOG) will meet in New Orleans, Louisiana, June 3–5, 2013 and in Phoenix, Arizona, October 7–9, 2013. For more information see: <http://nanog.org>

The *Internet Engineering Task Force* (IETF) will meet in Berlin, Germany, July 28–August 2, 2013 and in Vancouver, Canada, November 3–8, 2013. For more information see: <http://www.ietf.org/meeting/>

The *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT) will meet in Bangkok, Thailand, February 18–28, 2014. For more information see: <http://www.apricot.net>

## Call for Papers

*The Internet Protocol Journal* (IPJ) is published quarterly by Cisco Systems. The journal is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. The journal carries tutorial articles (“What is...?”), as well as implementation/operation articles (“How to...”). It provides readers with technology and standardization updates for all levels of the protocol stack and serves as a forum for discussion of all aspects of internetworking.

Topics include, but are not limited to:

- Access and infrastructure technologies such as: ISDN, Gigabit Ethernet, SONET, ATM, xDSL, cable, fiber optics, satellite, wireless, and dial systems
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, network computing, and Quality of Service
- Application and end-user issues such as: e-mail, Web authoring, server technologies and systems, electronic commerce, and application management
- Legal, policy, and regulatory topics such as: copyright, content control, content liability, settlement charges, “modem tax,” and trademark disputes in the context of internetworking

In addition to feature-length articles, IPJ contains standardization updates, overviews of leading and bleeding-edge technologies, book reviews, announcements, opinion columns, and letters to the Editor.

Cisco will pay a stipend of US\$1000 for published, feature-length articles. Author guidelines are available from Ole Jacobsen, the Editor and Publisher of IPJ, reachable via e-mail at [ole@cisco.com](mailto:ole@cisco.com)

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.



The Internet Protocol Journal, Cisco Systems  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA

ADDRESS SERVICE REQUESTED

PRSRT STD U.S. Postage PAID PERMIT No. 5187 SAN JOSE, CA
--

---

## The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

### Editorial Advisory Board

**Dr. Vint Cerf**, VP and Chief Internet Evangelist  
Google Inc, USA

**Dr. Jon Crowcroft**, Marconi Professor of Communications Systems  
University of Cambridge, England

**David Farber**  
Distinguished Career Professor of Computer Science and Public Policy  
Carnegie Mellon University, USA

**Peter Löthberg**, Network Architect  
Stupi AB, Sweden

**Dr. Jun Murai**, General Chair Person, WIDE Project  
Vice-President, Keio University  
Professor, Faculty of Environmental Information  
Keio University, Japan

**Dr. Deepinder Sidhu**, Professor, Computer Science &  
Electrical Engineering, University of Maryland, Baltimore County  
Director, Maryland Center for Telecommunications Research, USA

**Pindar Wong**, Chairman and President  
Verifi Limited, Hong Kong

*The Internet Protocol Journal is  
published quarterly by the  
Chief Technology Office,  
Cisco Systems, Inc.  
[www.cisco.com](http://www.cisco.com)  
Tel: +1 408 526-4000  
E-mail: [ipj@cisco.com](mailto:ipj@cisco.com)*

*Copyright © 2013 Cisco Systems, Inc.  
All rights reserved. Cisco, the Cisco  
logo, and Cisco Systems are  
trademarks or registered trademarks  
of Cisco Systems, Inc. and/or its  
affiliates in the United States and  
certain other countries. All other  
trademarks mentioned in this document  
or Website are the property of their  
respective owners.*

*Printed in the USA on recycled paper.*

