

The Internet Protocol Journal

December 2010

Volume 13, Number 4

*A Quarterly Technical Publication for
Internet and Intranet Professionals*

In This Issue

From the Editor	1
Emergency Services.....	2
Integrating Core BGP/MPLS Networks	18
Letter to the Editor	32
Book Review.....	33
Fragments	37

FROM THE EDITOR

I have recently started using both a smartphone and a tablet device for Internet access. Like millions of other Internet users, I have discovered the wonders of mobile applications that provide everything from the traditional Internet services (e-mail and web browsing) to specialized software that can pinpoint my location on a map, provide live currency-exchange calculations, give weather forecasts, and my favorite: play radio stations from all over the world. I am old enough to remember the orange glow from pre-transistor vacuum-tube radios, so having a customizable “world radio” in the form of an “app” on a smartphone seems almost like science fiction.

But radio is not the only traditional service that is now available over the Internet. Another prominent example is telephony or *Voice over IP* (VoIP). Not only is VoIP replacing traditional land lines in many places, the original circuit-switched telephone network is itself increasingly using VoIP technology in place of an infrastructure of land lines and dedicated switching equipment. An important aspect of traditional phone service is the notion of special numbers for *emergency services*. Such systems rely on a database of phone numbers and addresses that allow emergency personnel to dispatch responders to the correct location. This location identification becomes a lot more complicated if the caller is using an Internet-based calling service rather than a hard-wired telephone. The IETF has been tackling this problem in the *Emergency Context Resolution with Internet Technology* (ECRIT) working group. Our first article, by Hannes Tschofenig and Henning Schulzrinne, is an overview of the architecture this working group is developing.

According to the ITU-T, a *Next Generation Network* (NGN) is “...a packet-based network which can provide services including Telecommunication Services and is able to make use of multiple broadband, Quality of Service-enabled transport technologies in which service-related functions are independent from underlying transport-related technologies.” Paul Veitch, Paul Hitchen, and Martin Mitchell describe the integration of a standalone core BGP/MPLS VPN network into an NGN architecture.

Please check your subscription expiration date and renew online if you wish to continue receiving this journal. Click the “Subscriber Services” link at www.cisco.com/ipj to get to the login page. If you need any assistance just send e-mail to ipj@cisco.com and we will make the necessary changes for you.

—Ole J. Jacobsen, Editor and Publisher

ole@cisco.com

You can download IPJ
back issues and find
subscription information at:
www.cisco.com/ipj

ISSN 1944-1134

Emergency Services for Internet Multimedia

by Hannes Tschofenig, Nokia Siemens Networks and Henning Schulzrinne, Columbia University

Summoning the police, the fire department, or an ambulance in emergencies is one of the most important functions the telephone enables. As telephone functions move from circuit-switched to Internet telephony, telephone users rightfully expect that this core feature will continue to be available and work as well as it has in the past. Users also expect to be able to reach emergency assistance using new communication devices and applications, such as instant messaging or *Short Message Service* (SMS), and new media, such as video. In all cases, the basic objective is the same: The person seeking help needs to be connected with the most appropriate *Public Safety Answering Point* (PSAP), where call takers dispatch assistance to the caller's location. PSAPs are responsible for a particular geographic region, which can be as small as a single university campus or as large as a country.

The transition to Internet-based emergency services introduces two major structural challenges. First, whereas traditional emergency calling imposed no requirements on end systems and was regulated at the national level, Internet-based emergency calling needs global standards, particularly for end systems. In the old *Public Switched Telephone Network* (PSTN), each caller used a single entity, the landline or mobile carrier, to obtain services. For Internet multimedia services, network-level transport and applications can be separated, with the *Internet Service Provider* (ISP) providing IP connectivity service, and a *Voice Service Provider* (VSP) adding call routing and PSTN termination services. We ignore the potential separation between the Internet access provider, that is, a carrier that provides physical and data link layer network connectivity to its customers, and the ISP that provides network layer services. We use the term VSP for simplicity, instead of the more generic term *Application Server Provider* (ASP).

The documents that the IETF *Emergency Context Resolution with Internet Technology* (ECRIT) working group is developing support multimedia-based emergency services, and not just voice. As is explained in more detail later in this article, emergency calls need to be identified for special call routing and handling services, and they need to carry the location of the caller for routing and dispatch. Only the calling device can reliably recognize emergency calls, while only the ISP typically has access to the current geographical location of the calling device based on its point of attachment to the network. The reliable handling of emergency calls is further complicated by the wide variety of access technologies in use, such as *Virtual Private Networks* (VPNs), other forms of tunneling, firewalls, and *Network Address Translators* (NATs).

This article describes the architecture of emergency services as defined by the IETF and some of the intermediate steps as end systems and the call-handling infrastructure transition from the current circuit-switched and emergency-calling-unaware *Voice-over-IP* (VoIP) systems to a true any-media, any-device emergency calling system.

IETF Emergency Services Architecture

The emergency services architecture developed by the IETF ECRIT working group is described in [1] and can be summarized as follows: *Emergency calls are generally handled like regular multimedia calls, except for call routing.* The ECRIT architecture assumes that PSAPs are connected to an IP network and support the *Session Initiation Protocol* (SIP)^[2] for call setup and messaging. However, the calling user agent may use any call signaling or instant messaging protocol, which the VSP then translates into SIP.

Nonemergency calls are routed by a VSP, either to another subscriber of the VSP, typically through some SIP session border controller or proxy, or to a PSTN gateway. For emergency calls, the VSP keeps its call routing role, routing calls to the emergency service system to reach a PSAP instead. However, we also want to allow callers that do not subscribe to a VSP to reach a PSAP, using nothing but a standard SIP^[2] user agent (see [3] and [4] for a discussion about this topic); the same mechanisms described here apply. Because the Internet is global, it is possible that a caller's VSP resides in a regulatory jurisdiction other than where the caller and the PSAP are located. In such circumstances it may be desirable to exclude the VSP and provide a direct signaling path between the caller and the emergency network. This setup has the advantage of ensuring that all parties included in the call delivery process reside in the same regulatory jurisdiction.

As noted in the introduction, the architecture neither forces nor assumes any type of trust or business relationship between the ISP and the VSP carrying the emergency call. In particular, this design assumption affects how location is derived and transported.

Providing emergency services requires three crucial steps, which we describe in the following sections: recognizing an emergency call, determining the caller's location, and routing the call and location information to the appropriate emergency service system operating a PSAP.

Recognizing an Emergency Call

In the early days of PSTN-based emergency calling, callers would dial a local number for the fire or police department. It was recognized in the 1960s that trying to find this number in an emergency caused unacceptable delays; thus, most countries have been introducing single nationwide emergency numbers, such as 911 in North America, 999 in The United Kingdom, and 112 in all European Union countries.

This standardization became even more important as mobile devices started to supplant landline phones. In some countries, different types of emergency services, such as police or mountain rescue, are identified by separate numbers. Unfortunately, more than 60 different emergency numbers are used worldwide, many of which also have nonemergency uses in other countries, so simply storing the list of numbers in all devices is not feasible. In addition, hotels and university campuses often use dial prefixes, so an emergency caller in some European universities may actually have to dial 0112 to reach the fire department.

Because of this diversity, the ECRIT architecture decided to separate the concept of an emergency dial string, which remains the familiar and regionally defined emergency number, and a protocol identifier that is used for identifying emergency calls within the signaling system. The calling end system has to recognize the emergency (service) dial string and translate it into an emergency service identifier, which is an extensible set of *Uniform Resource Names* (URNs) defined in RFC 5031^[5]. A common example for such a URN, defined to reach the generic emergency service, is `urn:service.sos`. The emergency service URN is included in the signaling request as the destination and is used to identify the call as an emergency call. If the end system fails to recognize the emergency dial string, the VSP may also perform this service.

Because mobile devices may be sold and used worldwide, we want to avoid manually configuring emergency dial strings. In general, a device should recognize the emergency dial string familiar to the user and the dial strings customarily used in the currently visited country. The *Location-to-Service Translation Protocol* (LoST)^[6], described in more detail later, also delivers this information.

Some devices, such as smartphones, can define dedicated user interface elements that dial emergency services. However, such mechanisms must be carefully designed so that they are not accidentally triggered, for example, when the device is in a pocket.

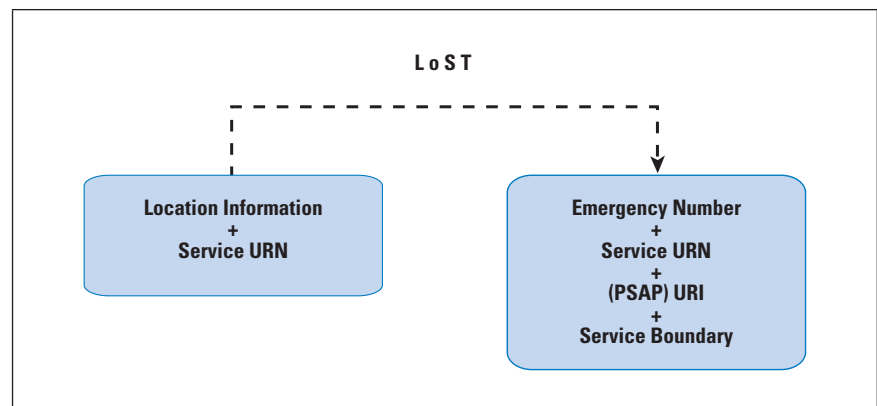
Emergency Call Routing

When an emergency call is recognized, the call needs to be routed to the appropriate PSAP. Each PSAP is responsible for only a limited geographic region, its service region, and some set of emergency services. For example, even in countries with a single general emergency number such as the United States, poison-control services maintain their own set of call centers. Because VSPs and end devices cannot keep a complete up-to-date mapping of all the service regions, a mapping protocol, LoST^[6], maps a location and service URN to a specific PSAP *Uniform Resource Identifier* (URI) and a service region.

LoST, illustrated in Figure 1, is a *Hypertext Transfer Protocol* (HTTP)-based query/response protocol where a client sends a request containing the location information and service URN to a server and receives a response containing the service URL, typically a SIP URL, the service region where the same information would be returned, and an indication of how long the information is valid. Both request and response are formatted as *Extensible Markup Language* (XML). For efficiency, responses are cached, because otherwise every small movement would trigger a new LoST request. As long as the client remains in the same service region, it does not need to consult the server again until the response returned reaches its expiration date. The response may also indicate that only a more generic emergency service is offered for this region. For example, a request for **urn:service:sos.marine** in Austria may be replaced by **urn:service:sos**. Finally, the response also indicates the emergency number and dial string for the respective service.

The number of PSAPs serving a country varies significantly. Sweden, for example, has 18 PSAPs, and the United States has approximately 6,200. Therefore, there is roughly one PSAP per 500,000 inhabitants in Sweden and one per 50,000 in the United States. As all-IP infrastructure is rolled out, smaller PSAPs may be consolidated into regional PSAPs. Routing may also take place in multiple stages, with the call being directed to an *Emergency Services Routing Proxy* (ESRP), which in turn routes the call to a PSAP, accounting for factors such as the number of available call takers or the language capabilities of the call takers.

Figure 1: High-Level Functions of Location-to-Service Translation (LoST) Protocol



Location Information

Emergency services need location information for three reasons: routing the call to the right PSAP, dispatching first responders (for example, policemen), and determining the right emergency service dial strings. It is clear that the location must be automatic for the first and third applications, but experience has shown that automated, highly accurate location information is vital to dispatching as well, rather than relying on callers to report their locations to the call taker.

Such information increases accuracy and avoids dispatch delays when callers are unable to provide location information because of language barriers, lack of familiarity with their surroundings, stress, or physical or mental impairment.

Location information for emergency purposes comes in two representations: geo(detic), that is, longitude and latitude, and civic, that is, street addresses similar to postal addresses. Particularly for indoor location, vertical information (floors) is very useful. Civic locations are most useful for fixed Internet access, including wireless hotspots, and are often preferable for specifying indoor locations, whereas geodetic location is frequently used for cell phones. However, with the advent of femto and pico cells, civic location is both possible and probably preferable because accurate geodetic information can be very hard to acquire indoors.

In almost all cases, location values are represented as *Presence Information Data Format Location Object* (PIDF-LO), an XML-based document to encapsulate civic and geodetic location information. The format of PIDF-LO is described in [7], with the civic location format updated in [8] and the geodetic location format profiled in [9]. The latter document uses the *Geography Markup Language* (GML) developed by the *Open Geospatial Consortium* (OGC) for describing commonly used location shapes.

Location can be conveyed either by value (“LbyV”) or by reference (“LbyR”). For the former, the XML location object is added as a message body in the SIP message. Location by value is particularly appropriate if the end system has access to the location information; for example, if it contains a *Global Positioning System* (GPS) receiver or uses one of the location configuration mechanisms described later in this section. In environments where the end host location changes frequently, the LbyR mechanism might be more appropriate. In this case, the LbyR is an HTTP/*Secure HTTP* (HTTPS) or SIP/*Secure SIP* (SIPS) URI, which the recipient needs to resolve to obtain the current location. Terminology and requirements for the LbyR mechanism are available in [10].

An LbyV and an LbyR can be obtained through location configuration protocols, such as the *HTTP Enabled Location Delivery* (HELD) protocol^[11] or *Dynamic Host Configuration Protocol* (DHCP)^[12, 13]. When obtained, location information is required for LoST queries, and that information is added to SIP messages^[14].

The requirements for location accuracy differ between routing and dispatch. For call routing, city or even county-level accuracy is often sufficient, depending on how large the PSAP service areas are, whereas first responders benefit greatly when they can pinpoint the caller to a particular building or, better yet, apartment or office for indoor locations, and an outdoor area of at most a few hundred meters. This detailed location information avoids having to search multiple buildings, for example, for medical emergencies.

As mentioned previously, the ISP is the source of the most accurate and dependable location information, except for cases where the calling device has built-in location capabilities, such as GPS, when it may have more accurate location information. For landline Internet connections such as DSL, cable, or fiber-to-the-home, the ISP knows the provisioned location for the network termination, for example. The IETF GEOPRIV working group has developed protocol mechanisms, called *Location Configuration Protocols*, so that the end host can request and receive location information from the ISP. The Best Current Practice document for emergency calling^[15] enumerates three options that clients should universally support: DHCP civic^[16] and geo^[12] (with a revision of RFC 3825 in progress^[17]), and HELD^[11]. HELD uses XML query and response objects carried in HTTP exchanges. DHCP does not use the PIDF-LO format, but rather more compact binary representations of locations that require the endpoint to construct the PIDF-LO.

Particularly for cases where end systems are not location-capable, a VSP may need to obtain location information on behalf of the end host^[18].

Obtaining at least approximate location information at the time of the call is time-critical, because the LoST query can be initiated only after the calling device or VSP has obtained location information. Also, to accelerate response, it is desirable to transmit this location information with the initial call signaling message. In some cases, however, location information at call setup time is imprecise. For example, a mobile device typically needs 15 to 20 seconds to get an accurate GPS location “fix,” and the initial location report is based on the cell tower and sector. For such calls, the PSAP should be able to request more accurate location information either from the mobile device directly or the *Location Information Server* (LIS) operated by the ISP. The SIP event notification extension, defined in RFC 3265^[19], is one such mechanism that allows a PSAP to obtain the location from an LIS. To ensure that the PSAP is informed only of pertinent location changes and that the number of notifications is kept to a minimum, event filters^[20] can be used.

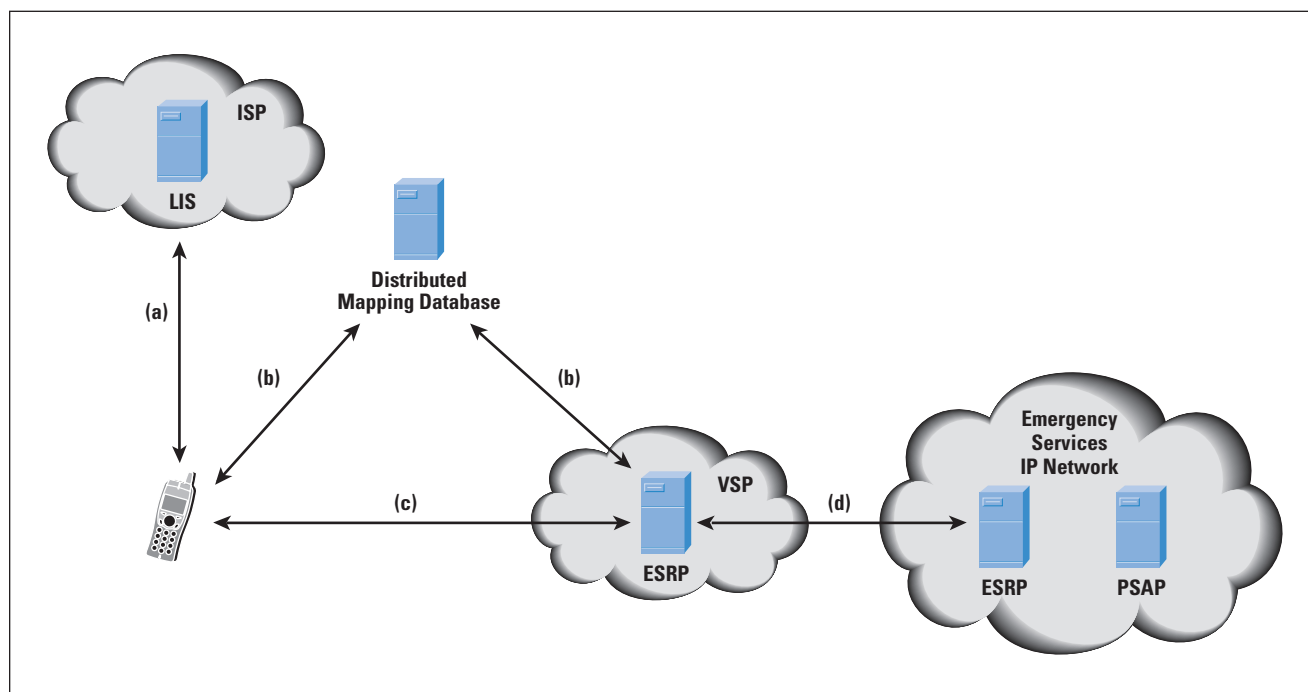
The two-stage location refinement mechanism described previously works best when location is provided by reference (LbyR) in the SIP INVITE call setup request. The PSAP subscribes to the LbyR provided in the SIP exchange and the LbyR refers to the LIS in the ISP’s network. In addition to a SIP URI, the LbyR message can also contain an HTTP/HTTPS URI. When such a URI is provided, an HTTP-based protocol can be used to retrieve the current location^[21].

Obligations

This section discusses the requirements the different entities need to satisfy, based on Figure 2. A more detailed description can be found in [15].

Note that this narration focuses on the final stage of deployment and does not discuss the transition architecture, in which some implementation responsibilities can be rearranged, with an effect on the overall functions offered by the emergency services architecture. A few variations were introduced to handle the transition from the current system to a fully developed ECRIT architecture.

Figure 2: Main Components Involved in an Emergency Call



With the work on the IETF emergency architecture, we have tried to balance the responsibilities among the participants, as described in the following sections.

End Hosts

An end host, through its VoIP application, has three main responsibilities: it has to attempt to obtain its own location, determine the URI of the appropriate PSAP for that location, and recognize when the user places an emergency call by examining the dial string. The end host operating system may assist in determining the device location.

The protocol interaction for location configuration is indicated as interface (a) in Figure 2; numerous location configuration protocols have been developed to provide this capability.

A VoIP application needs to support the LoST protocol^[6] in order to determine the emergency service dial strings and the PSAP URI. Additionally, the device needs to understand the service identifiers, defined in [5].

As currently defined, it is assumed that SIP can reach PSAPs, but PSAPs may support other signaling protocols, either directly or through a protocol translation gateway. The LoST retrieval results indicate whether other signaling protocols are supported. To provide support for multimedia, use of different types of codecs may be required; details are available in [15].

ISP

The ISP has to make location information available to the endpoint through one or more of the location configuration protocols.

In order to route an emergency call correctly to a PSAP, an ISP may initially disclose the approximate location for routing to the endpoint and give more precise location information later, when the PSAP operator dispatches emergency personnel. The functions required by the IETF emergency services architecture are restricted to the disclosure of a relatively small amount of location information, as discussed in [22] and in [23].

The ISP may also operate a (caching) LoST server to improve the robustness and reliability of the architecture. This server lowers the round-trip time for contacting a LoST server, and the caches are most likely to hold the mappings of the area where the emergency caller is currently located.

When ISPs allow Internet traffic to traverse their network, the signaling and media protocols used for emergency calls function without problems. Today, there are no legal requirements to offer prioritization of emergency calls over IP-based networks. Although the standardization community has developed a range of *Quality of Service* (QoS) signaling protocols, they have not experienced widespread deployment.

VSP

SIP does not mandate that call setup requests traverse SIP proxies; that is, SIP messages can be sent directly to the user agent. Thus, even for emergency services it is possible to use SIP without the involvement of a VSP. However, in terms of deployment, it is highly likely that a VSP will be used. If a caller uses a VSP, this VSP often forces all calls, emergency or not, to traverse an outbound proxy or *Session Border Controller* (SBC) operated by the VSP. If some end devices are unable to perform a LoST lookup, VSP can provide the necessary functions as a backup solution.

If the VSP uses a signaling or media protocol that the PSAP does not support, it needs to translate the signaling or media flows.

VSPs can assist the PSAP by providing identity assurance for emergency calls; for example, using [30], thus helping to prosecute prank callers. However, the link between the subscriber information and the real-world person making the call is weak.

In many cases, VSPs have, at best, only the credit card data for their customers, and some of these customers may use gift cards or other anonymous means of payment.

PSAP

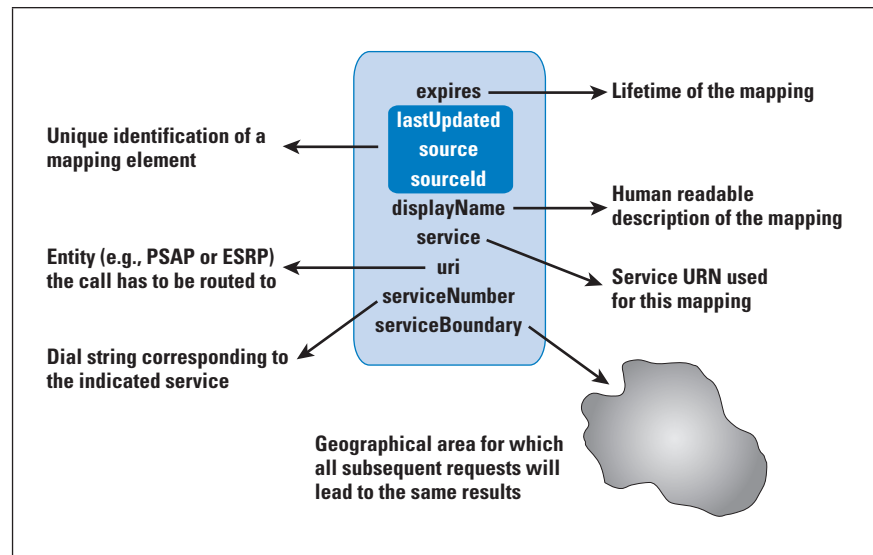
The emergency services Best Current Practice document [15] discusses only the standardization of the interfaces from the VSP and ISP toward PSAPs and some parts of the PSAP-to-PSAP call transfer mechanisms that are necessary for emergency calls to be processed by the PSAP. Many aspects related to the internal communication within a PSAP, between PSAPs as well as between a PSAP and first responders, are beyond the scope of the IETF specification.

When emergency calling has been fully converted to Internet protocols, PSAPs must accept calls from any VSP, as shown in interface (d) of Figure 2. Because calls may come from all sources, PSAPs must develop mechanisms to reduce the number of malicious calls, particularly calls containing intentionally false location information. Assuring the reliability of location information remains challenging, particularly as more and more devices are equipped with *Global Navigation Satellite Systems* (GNSS) receivers, including GPS and Galileo, allowing them to determine their own location^[24]. However, it may be possible in some cases to check the veracity of the location information an endpoint provides by comparing it against infrastructure-provided location information; for example, a LIS-determined location.

Mapping Architecture

So far we have described LoST as a client-server protocol. Similar to the *Domain Name System* (DNS), a single LoST server does not store the mapping elements for all PSAPs worldwide, for both technical and administrative reasons. Thus, there is a need to let LoST servers interact with other LoST servers, each covering a specific geographical region. Working together, LoST servers form a distributed mapping database, with each server carrying mapping elements, as shown in Figure 3. LoST servers may be operated by different entities, including the ISP, the VSP, or another independent entity, such as a governmental agency. Typically, individual LoST servers offer the necessary mapping elements for their geographic regions to others. However, LoST servers may also cache mapping elements of other LoST servers either through data synchronization mechanisms (for example, FTP or exports from a *Geographical Information System* [GIS] or through a specialized protocol^[25]) or by regular usage of LoST. This caching improves performance and increases the robustness of the system.

Figure 3: Mapping Element



A detailed description of the mapping architecture with examples is available in [29].

Steps Toward an IETF Emergency Services Architecture

The architecture described so far requires changes both in already-deployed VoIP end systems and in the existing PSAPs. The speed of transition and the path taken vary between different countries, depending on funding and business incentives. Therefore, it is generally difficult to argue whether upgrading endpoints or replacing the emergency service infrastructure will be easier. In any case, the transition approaches being investigated consider both directions. We can distinguish roughly four stages of transition (Note: The following descriptions omit many of the details because of space constraints):


1. Initially, VoIP end systems cannot place emergency calls at all; for example, many software clients, such as *GoogleTalk*, cannot place emergency calls.
2. In a second stage, VoIP callers manually configure their location, and emergency calls are routed to the appropriate PSAP as circuit-switched calls through PSTN gateways using technologies similar to mobile calls. This level of service is now offered in some countries for PSTN-replacement VoIP services; that is, VoIP services that are offered as replacement for the home phone. In the United States, this service is known as the “NENA I2” service.
3. In a third stage, PSAPs maintain two separate infrastructures, one for calls arriving through an IP network and the traditional infrastructure.
4. In the final stage, all calls, including those from traditional cell phones and analog landline phones, reach the PSAP through IP networks, with the traditional calls converted to the ECRIT requirements by the carriers or the emergency service infrastructure.

If devices are used in environments without location services, the VSP's SIP proxy may need to insert location information based on estimates or subscriber data. These cases are described briefly in the following sections.

Traditional Endpoints

Figure 4 shows an emergency services architecture with traditional endpoints. When the emergency caller dials the Europeanwide emergency number 112 (step 0), the device treats it as any other call without recognizing it as an emergency call; that is, the dial string provided by the endpoint that may conform to RFC 4967^[26] or RFC 3966^[27] is signaled to the VSP (step 1). Recognition of the dial string is then left to the VSP for processing or sorting; the same is true for location retrieval (step 2) and routing to the nearest (or appropriate) PSAP (step 3). Dial-string recognition, location determination, and call routing are simpler to carry out using a fixed device and the voice and application service provided through the ISP than they are when the VSP and the ISP are two separate entities.

Figure 4: Emergency Services Architecture with Traditional Endpoints



There are two main challenges to overcome when dealing with traditional devices: First, the VSP must discover the LIS that knows the location of the IP-based end host. The VSP is likely to know only the IP address of that device, visible in the call signaling that arrives at the VSP. When a LIS is discovered and contacted and some amount of location information is available, then the second challenge arises, namely, how to route the emergency call to the appropriate PSAP. To accomplish the latter task it is necessary to have some information about the PSAP boundaries available.

Reference [15] does not describe a complete and detailed solution but uses building blocks specified in ECRIT. Still, this deployment scenario shows many constraints:

- Only the emergency numbers configured at the VSP are understood. This situation may lead to cases where a dialed emergency number is not recognized.
- Using the IP address to find the ISP is challenging and may, in case of mobility protocols and VPNs, lead to wrong results.
- Security concerns might arise when a potentially large number of VSPs or ASPs are able to retrieve location information from an ISP. It is likely that only authorized VSP and ASPs will be granted access. Hence, it is unlikely that such a solution would work smoothly across national boundaries.
- When the user agent does not recognize the emergency call, functions such as call waiting, call transfer, three-way call, flash hold, and outbound call blocking cannot be disabled.
- The user-agent software may block callbacks from the PSAP.
- Privacy settings may not get considered and identity may get disclosed to unauthorized parties. These identity privacy features exist in some jurisdictions even in emergency situations.
- Certain VoIP call features may not be supported, such as REFER (for conference call and transfer to secondary PSAP) and *Globally Routable UA URI* (GRUU).
- User agents will not convey location information to the VSP (even if available).

Partially Upgraded End Hosts

A giant step forward in simplifying the handling of IP-based emergency calls is to provide the end host with some information about the ISP so that LIS discovery is possible. The end host may, for example, learn the ISP's domain name by using LIS discovery^[28], or might even obtain a *Location by Reference* (LbyR) through the DHCP-URI option^[13] or through HELD^[11]. The VSP can then either resolve the LbyR in order to route the call or use the domain to discover a LIS using DNS.

Additional software upgrades at the end device may allow for recognition of emergency calls based on some preconfigured emergency numbers (for example, 112 and 911) and allow for the implementation of other emergency service-related features, such as disabling silence suppression during emergency calls.

Outlook

In most countries, national and sometimes regional telecommunications regulators, such as the *Federal Communications Commission* (FCC) and individual states, or the European Union, strongly influence how emergency services are provided, who pays for them, and the obligations that the various parties have. Regulation is, however, still at an early stage: in most countries current requirements demand only manual update of location information by the VoIP user. The ability to obtain location information automatically is, however, crucial for reliable emergency service operation, and it is required for nomadic and mobile devices. (Nomadic devices remain in one place during a communication session, but are moved frequently from place to place. Laptops with Wi-Fi interfaces are currently the most common nomadic devices.)

Regulators have traditionally focused on the national or, at most, the European level, and the international nature of the Internet poses new challenges. For example, mobile devices are now routinely used beyond their country of purchase and, unlike traditional cellular phones, need to support emergency calling functions. It appears likely that different countries will deploy IP-based emergency services over different time horizons, so travelers may be surprised to find that they cannot call for emergency assistance outside their home country.

The separation between Internet access and application providers on the Internet is one of the most important differences to existing circuit-switched telephony networks. A side effect of this separation is the increased speed of innovation at the application layer, and the number of new communication mechanisms is steadily increasing. Many emergency service organizations have recognized this trend and advocated for the use of new communication mechanisms, including video, real-time text, and instant messaging, to offer improved emergency calling support for citizens. Again, this situation requires regulators to rethink the distribution of responsibilities, funding, and liability.

Many communication systems used today lack accountability; that is, it is difficult or impossible to trace malicious activities back to the persons who caused them. This problem is not new, because pay phones and prepaid cell phones have long offered mischief makers the opportunity to place hoax calls, but the weak user registration procedures, the lack of deployed end-to-end identity mechanisms, and the ease of providing fake location information increases the attack surface at PSAPs. Attackers also have become more sophisticated over time, and Botnets that generate a large volume of automated emergency calls to exhaust PSAP resources, including call takers and first responders, are not science fiction.

References

- [1] Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, "Framework for Emergency Calling Using Internet Multimedia," Internet Draft, work in progress, **draft-ietf-ecrit-framework-11**, July 2010.
- [2] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol," RFC 3261, June 2002.
- [3] Winterbottom, J., Thomson, M., Tschofenig, H., and H. Schulzrinne, "ECRIT Direct Emergency Calling," Internet Draft, work in progress, **draft-winterbottom-ecrit-direct-02.txt**, March 2010.
- [4] Schulzrinne, H., McCann, S., Bajko, G., Tschofenig, H., and D. Kroesenberg, "Extensions to the Emergency Services Architecture for Dealing with Unauthenticated and Unauthorized Devices," Internet Draft, work in progress, **draft-ietf-ecrit-unauthenticated-access-00.txt**, September 2010.
- [5] Schulzrinne, H., "A Uniform Resource Name (URN) for Emergency and Other Well-Known Services," RFC 5031, January 2008.
- [6] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol," RFC 5222, August 2008.
- [7] Peterson, J., "A Presence-based GEOPRIV Location Object Format," RFC 4119, December 2005.
- [8] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)," RFC 5139, February 2008.
- [9] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations," RFC 5491, March 2009.
- [10] R. Marshall, "Requirements for a Location-by-Reference Mechanism," RFC 5808, May 2010.
- [11] M. Barnes, "HTTP Enabled Location Delivery (HELD)," RFC 5985, September 2010.
- [12] Polk, J., Schnizlein, J., and M. Linsner, "Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information," RFC 3825, July 2004.

- [13] Polk, J., “Dynamic Host Configuration Protocol (DHCP) IPv4 and IPv6 Option for a Location Uniform Resource Identifier (URI),” Internet Draft, work in progress, **draft-ietf-geopriv-dhcp-lbyr-uri-option-08**, July 2010.
- [14] Polk, J., Rosen, B., and J. Peterson, “Location Conveyance for the Session Initiation Protocol,” Internet Draft, work in progress, **draft-ietf-sipcore-location-conveyance-03**, July 2010.
- [15] Rosen, B. and J. Polk, “Best Current Practice for Communications Services in Support of Emergency Calling,” Internet Draft, work in progress, **draft-ietf-ecrit-phonebcp-15**, July 2010.
- [16] Schulzrinne, H., “Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information,” RFC 4776, November 2006.
- [17] Polk, J., Schnizlein, J., Linsner, M., and B. Aboba, “Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information,” Internet Draft, work in progress, **draft-ietf-geopriv-rfc3825bis-11**, July 2010.
- [18] Winterbottom, J., Thomson, M., Tschofenig, H., and R. Barnes, “Use of Device Identity in HTTP-Enabled Location Delivery (HELD),” Internet Draft, work in progress, **draft-ietf-geopriv-held-identity-extensions-04**, June 2010.
- [19] Roach, A., “Session Initiation Protocol (SIP)-Specific Event Notification,” RFC 3265, June 2002.
- [20] Mahy, R., Rosen, B., and H. Tschofenig, “Filtering Location Notifications in the Session Initiation Protocol,” Internet Draft, work in progress, **draft-ietf-geopriv-loc-filters-11**, March 2010.
- [21] Winterbottom, J., Tschofenig, H., Schulzrinne, H., Thomson, M., and M. Dawson, “A Location Dereferencing Protocol Using HELD,” Internet Draft, work in progress, **draft-ietf-geopriv-deref-protocol-01**, September 2010.
- [22] Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, “Location Hiding: Problem Statement and Requirements,” Internet Draft, work in progress, **draft-ietf-ecrit-location-hiding-req-04**, Feb 2010.
- [23] Barnes, R., and M. Lepinski, “Using Imprecise Location for Emergency Context Resolution,” Internet Draft, work in progress, **draft-ietf-ecrit-rough-loc-03**, August 2010.

- [24] Tschofenig, H., Schulzrinne, H., and B. Aboba, "Trustworthy Location Information," Internet Draft, work in progress, **draft-tschofenig-ecrit-trustworthy-location-00**, September 2010.
- [25] Schulzrinne, H., and H. Tschofenig, "Synchronizing Location-to-Service Translation (LoST) Servers," Internet Draft, work in progress, **draft-ietf-ecrit-lost-sync-10**, March 2010.
- [26] B. Rosen, "Dial String Parameter for the Session Initiation Protocol Uniform Resource Identifier," RFC 4967, July 2007.
- [27] H. Schulzrinne "The tel URI for Telephone Numbers," RFC 3966, December 2004.
- [28] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)," RFC 5986, September 2010.
- [29] H. Schulzrinne, "Location-to-URL Mapping Architecture and Framework," RFC 5582, September 2009.
- [30] C. Jennings, J. Peterson, and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks," RFC 3325, November 2002.

HENNING SCHULZRINNE, Levi Professor of Computer Science at Columbia University, received his Ph.D. from the University of Massachusetts in Amherst, Massachusetts. He was an MTS at AT&T Bell Laboratories and an associate department head at GMD-Fokus (Berlin) before joining the Computer Science and Electrical Engineering departments at Columbia University. He served as chair of Computer Science from 2004 to 2009. Protocols that he co-developed, such as RTP, RTSP, and SIP, are now Internet standards, used by almost all Internet telephony and multimedia applications. His research interests include Internet multimedia systems, ubiquitous computing, and mobile systems. He is a Fellow of the IEEE. E-mail: hgs@cs.columbia.edu

HANNES TSCHOFENIG received a Diploma degree from the University of Klagenfurt, Austria. He joined Siemens Corporate Technology, Munich, in 2001 and joined Nokia Siemens Networks in April 2007 to move to Finland in December 2007, where he focuses on standards development. Most of his time is dedicated to the participation in the Internet Engineering Task Force (IETF) where he, among other responsibilities, co-chaired the ECRIT working group from 2005 to early 2010. Additionally, he co-chairs the Next Generation 112 Technical Committee of the European Emergency Number Association (EENA) and contributes to the technical specifications developed within the National Emergency Number Association (NENA), and he co-organized the SDO emergency services workshop series. In March 2010 he joined the Internet Architecture Board (IAB). E-mail: hannes.tschofenig@nsn.com

Integration of Core BGP/MPLS VPN Networks

by Paul Veitch, Paul Hitchen, and Martin Mitchell, BT Innovate & Design

This article explores the architectural and operational challenges involved in integrating an existing standalone core *Border Gateway Protocol* (BGP)/*Multiprotocol Label Switching* (MPLS) VPN network onto a target *Next-Generation Network* (NGN). The rationale for consolidating and transforming multiple networks is explained, mainly in terms of potential cost savings and operational simplification achieved by the network operator. The article specifically focuses on the MPLS *Carrier-supporting-Carrier* (CsC) architectural framework, which allows the serving nodes of one MPLS VPN network to be interconnected through the serving nodes of another MPLS VPN network. The required architectural building blocks to implement CsC, the manner in which routing protocols must interact, as well as end-to-end packet flow and label encapsulation are all explained. The main design and operational challenges, including maintaining performance levels for customers, network resiliency, fault-handling, and capacity management, are also addressed in this article.

Network operators are under increasing pressure to deliver exceptional levels of customer experience and service while decreasing the capital and operational cost base of their networks. Many operators have traditionally built multiple network platforms, each of which has been uniquely designed to meet the requirements of specific services targeted at specific customer markets, such as voice, broadband IP, *Virtual Private Networks* (VPNs), etc.

In a bid to remain competitive and achieve cost reductions and operational simplifications, many operators have built all IP-based NGNs. The principal transformational benefits of an NGN with a single protocol such as IP at its heart include versatility in catering for multiple traffic requirements (for example, by employing IP *Quality-of-Service* [QoS] techniques), the ability to introduce novel and reusable services and features in a flexible manner, and the potential to maximise vendor interworking due to standards-based technology.

When a network operator builds an NGN, the challenge remains as to how to migrate *existing* networks and customers onto the new platform. The full commercial benefits of an NGN can be properly realised only after legacy networks are either consolidated or phased out completely. Many important factors must be considered, including the cost benefits, the potential effect on end customers, and the operational approach to carrying out migrations. These concerns must be weighed against the commercial and business risks associated with the alternative approach of sustaining and running multiple standalone platforms indefinitely.

This article focuses on a specific scenario: how to integrate an existing BGP/MPLS VPN network that provides VPN services to a corporate customer base with a “target” NGN. Following a brief overview of MPLS VPN services and networks, the rationale for consolidating multiple MPLS VPN networks is explained, mainly in terms of potential cost savings and operational simplification achieved by the network operator. The article then details the MPLS CsC architectural framework that allows the serving nodes or *Points of Presence* (POPs) of one MPLS VPN network to be interconnected to the serving nodes of another MPLS VPN network. The way in which routing protocols must interact and the subsequent effect on end-to-end packet forwarding across a CsC-enabled core network are explained. The principal design and operational challenges introduced by integrating core MPLS networks are then outlined, including maintaining performance levels, network resiliency, fault management, and capacity management.

The Business Case for MPLS VPN Network Consolidation

VPNs are an attractive solution to serve the enterprise networking requirements of a wide range of businesses from *Small-to-Medium Enterprises* (SMEs) to multinational “blue-chip” corporate organisations. Essentially, VPNs provide a transparent network infrastructure that allows multiple customer sites to communicate over a shared backbone network, as though they are using their own private network, regardless of geographical location. Typical applications that run across an organisation’s VPN include corporate Intranet, mail services, and *Voice-over-IP* (VoIP) telephony.


Although distinct categories of VPN networking technology exist^[1], this article focuses exclusively on “Layer 3” BGP/MPLS VPNs, as defined in RFC 4364^[2] and other related Internet Drafts. Such networks have been deployed for more than 10 years and have seen significant growth during that period.

The critical core network elements of a provider-provisioned BGP/MPLS VPN network are *Provider Edge* (PE) and *Provider Core* (P) routers, as shown in Figure 1.

PE routers terminate customer access circuits, whereas P routers perform packet forwarding and typically do not have directly connected customer access circuits. PE routers perform label encapsulation and de-encapsulation, P routers run label switching, and both operate control-plane protocols that build MPLS *Label Switched Paths* (LSPs) from each PE to each other PE. Many protocols can be used to establish these LSPs; a commonly deployed approach uses the *Label Distribution Protocol* (LDP) in conjunction with an *Interior Gateway Protocol* (IGP), such as *Open Shortest Path First* (OSPF).

When a PE forwards a VPN-addressed packet across the core, it adds an inner MPLS label to identify the VPN of which the packet is a member and then an outer MPLS label to identify the egress PE router. Any intermediate P routers switch the packet to the egress PE using the outer label only. The egress PE uses the inner label to determine which VPN or port to forward the packet to.

Figure 1: Overview of BGP/MPLS VPN Network



The *Customer Edge* (CE) router is not considered part of the provider's core network. It acts as a peer of the PE router, but not a peer of other CE routers. Each PE router supports multiple routing and forwarding tables, called *Virtual Route Forwarding* (VRF) tables. VRF routes are logically separate, and they may contain IP prefixes received from the CE router that overlap with addresses in other VRFs. (For example, in Figure 1, VPN_A, site 1 has the same private routes as VPN_B, site 3.) VPNs are formed by defining individual customer accesses to be members of a specific VRF table, with several sites formed on one PE by defining all sites to use the same VRF table or allocating each site a VRF table and controlling connectivity through selective import and export of the IP routes of each VRF table.

The PE routers use an extended variant of BGP for signaling between themselves and propagating information about the actual routes of each VPN, as well as the inner MPLS label. The extended BGP, referred to as *Multiprotocol BGP*, carries each VPN route together with two new fields, the *Route Distinguisher* (RD) and the *Route Target* (RT), a form of extended BGP Community.

The RD is added to each VPN route to ensure that routes from different customers are unique; BGP treats VPN routes as equal only if both the RD and the IP prefix mask are equal. BGP uses RTs to indicate a group of routes, thus defining VPN membership information for exchange between PEs.

Maintenance Costs of BGP/MPLS VPN Networks

As detailed in the previous section, the main core components of a VPN network based on BGP/MPLS technology are the PE and P routers. Although not shown in detail in Figure 1, another critical element of a core VPN network is the *Wide-Area Network* (WAN) topology that interconnects the P (core) routers residing in specific service nodes, also called POPs. The WAN topology is essentially the way in which transmission links—typically *Synchronous Optical Network* (SONET)/*Packet over SONET/SDH* (PoS), Gigabit Ethernet, or 10 Gigabit Ethernet—are used to interconnect the POPs together.

It follows that maintenance costs associated with a self-contained MPLS VPN network will be incurred for PE and P routers, as well as the interconnecting WAN transmission links. These maintenance costs will split into capital and operational elements.

Capital expenditures are required on an ongoing basis for all IP router infrastructure (PE and P routers), for example, to upgrade hardware to meet increasing capacity demands, replace faulty line cards and processors, or replace end-of-life hardware with newer equipment. Capital expenditures are also needed on WAN links, for example, to replace faulty line cards and optics, as well as to deploy increased capacity transmission links to cater for traffic growth across the core network. Further capital costs accrue from accommodation-related aspects such as power, racking, and air conditioning.

Additional maintenance costs reside in the operational space. For example, if an MPLS VPN network has 40 POP locations, each with a pair of P (core) routers, the 80 core routers will consume a certain amount of operational team resources for critical maintenance, scheduled maintenance activities, and ongoing monitoring and reporting of router status (processors and line cards).

Benefits of Core Integration

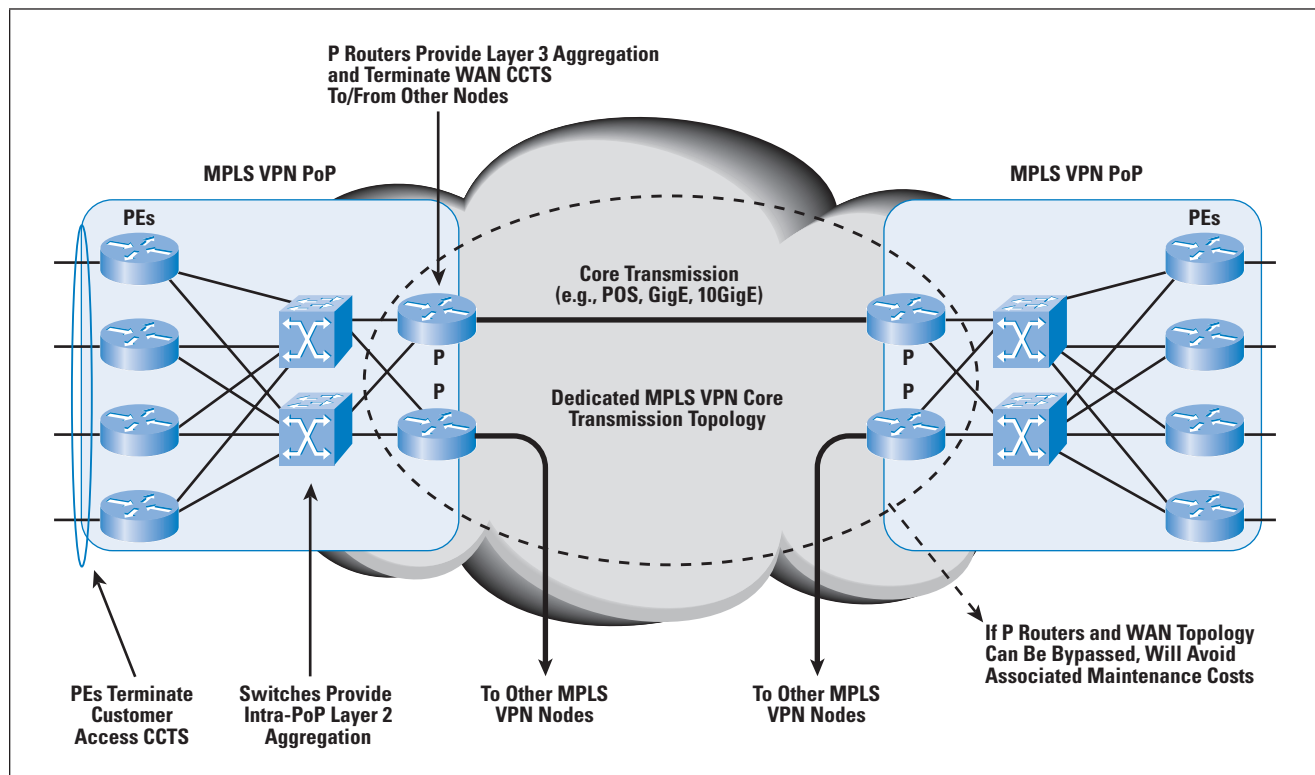
If a network operator has deployed an IP-based NGN alongside an existing MPLS VPN network, the question should be asked: can the existing MPLS VPN network be integrated onto the NGN so as to avoid some or all of the previously stated maintenance costs? One approach would be to target the P (core) routers and WAN transmission links for eventual removal (Figure 2) and replacement by suitable connectivity of the MPLS VPN nodes to the NGN network. The VPN PE routers that often terminate large volumes of customer access circuits and host the rich service-related functions for corporate VPN services can essentially be left *in situ*, minimising the effect on end customers and confining the integration of networks to the inner part of the core infrastructure. The way in which this goal can actually be achieved in practice is detailed in the next section.

The main benefits that can be accrued for the network operator are as follows:

- Substantial cost avoidance for maintaining and upgrading P (core) routers and dedicated WAN links for the existing MPLS VPN network can be achieved (Figure 2). As much as a 35-percent reduction of fixed inner core capital costs is possible.
- If the technical solution for core integration can be made as reusable as possible, then in addition to allowing integration of “same provider” core networks, the network operator could provide the capability on a wholesale basis for other service providers. This capability could be a potentially significant source of new revenue.
- From an operational perspective, integration of core networks should lend itself to a singular and much more streamlined approach to capacity planning, fault management, and network monitoring.

The combination of all these benefits can produce a compelling business case for network operators to consolidate core MPLS-based network platforms.

Figure 2: MPLS VPN Network Showing Inner Core Components Targeted for Replacement



Carrier-supporting-Carrier Framework

Carrier-supporting-Carrier (CsC) is a term used to describe a situation where one network, designated the *customer carrier*, is permitted to use a segment of another network, designated the *backbone carrier*^[3]. Although the term “Carrier of Carriers” is also used to describe the same architectural framework, this article uses Carrier-supporting-Carrier for consistency. In principle, the two “carrier” networks could belong to the same organisation, or could belong to two different organisations. Whatever the case, there is no reason why the backbone carrier cannot support multiple customer carrier networks. Furthermore, the customer carrier network itself can be either a BGP/MPLS VPN network providing Layer 3 VPN services or an *Internet Service Provider* (ISP) network^[3].

A network operator with an existing BGP/MPLS VPN network infrastructure that has also built an IP-based NGN based on BGP/MPLS technology as per RFC 4364^[2] could choose to exploit the CsC architectural framework to merge the two core networks. In such a scenario, the existing BGP/MPLS VPN network that serves the needs of VPN business customers would be viewed as the “customer carrier,” whereas the NGN network would be positioned as the “backbone carrier.”

Physical Connectivity and CsC VRF Creation

In order to integrate an existing BGP/MPLS VPN network such as that shown in Figure 2, with an NGN core belonging to the same or different organisation, the NGN network must be enabled to act as a backbone carrier. Assuming the NGN network is configured to support BGP/MPLS VPNs as per RFC 4364^[2], it comprises PE and P router core infrastructure. The PE routers of the NGN acting as the backbone carrier are denoted “CsC-PEs.” The PE routers of the existing BGP/MPLS VPN network, that is, the customer carrier network that is being itself integrated with the NGN core, are denoted “CsC-CEs.”

As shown in Figure 3, the NGN backbone carrier network provides MPLS VPN service to the customer carrier network using its own VRF table enabled on the CsC-PE. One important distinction between normal MPLS VPN service and CsC is the fact that traffic passed between the CsC-CE and CsC-PE is labeled rather than native IP^[3, 4].

The CsC architecture is designed such that the backbone carrier network—the network provider’s NGN network—needs to know only about internal routes within the customer carrier network. This setup allows formation of full “any-to-any” logical connectivity between the customer carrier routers, which in this scenario are the PE routers of the existing BGP/MPLS VPN network providing VPN services to end customers.

Furthermore, the backbone carrier routers themselves do not need to retain route prefix information for the end-customer VPNs connected to the customer carrier network because the end-customer traffic is transported over a second level of VRF tables that bear relevance only to the customer carrier itself, that is, the endpoint CsC-CEs. This *nesting* of MPLS VPN networks emphasises the inherent scalability of the CsC architecture. The CsC backbone carrier is effectively behaving like “proxy” P routers for the customer carrier network.

Figure 3: MPLS VPN “Customer Carrier” Network Connected Across NGN “Backbone Carrier”




Figure 3 also shows the physical connectivity between the customer carrier network and backbone carrier NGN. Because many large-scale BGP/MPLS network deployments comprise large numbers of PE devices in the same service node or POP, there is often a Layer 2 Ethernet switch acting as an “intra-POP” aggregator. It is convenient to allow physical connectivity between the BGP/MPLS VPN service node and the CsC-PE in the NGN network using this aggregation switch. One or more *Virtual LANs* (VLANs) can be configured across this physical trunk to provide logical Layer 2 connectivity into the CsC-PE on the NGN, and be associated with the CsC VRF on that device. The Layer 2 switch also provides direct intra-POP connectivity between CsC-CEs present on the same VLANs.

Control-Plane Routing Protocols


The previous section described the physical connectivity between BGP/MPLS VPN service nodes and the target NGN, with creation of a specific VRF route on the CsC-PEs. This section addresses the way in which the internal routes of the CsC-CEs (that is, the PE routers belonging to the customer carrier BGP/MPLS VPN network) are advertised into this VRF table.

Optional routing protocols include the use of an IGP such as OSPF, or *Exterior Gateway Protocols* (EGPs) such as BGP. With an IGP like OSPF^[5], the routing protocol itself is used for route exchange between the CsC-CEs and CsC-PEs, and must be used in conjunction with an LDP^[6] for MPLS label exchange between the CsC-CEs and CsC-PEs.

Separating the IP prefix and label allocation protocols between an IGP and LDP can introduce complexities with potential divergence between the two control planes. Such divergence in the extreme case can lead to partial or complete loss in forwarding. Use of an EGP like BGP, however, can be used to implement CsC as a single IP prefix and Label Allocation control-plane protocol between CsC-CE and CsC-PE. Piggybacking MPLS label-mapping information in the BGP update messages helps ensure that an IP prefix and its associated MPLS label are always synchronised in their delivery. The way in which this synchronisation is achieved is documented in RFC 3107^[7]. BGP has the benefit of being a mature protocol for use either within the same network organisation or between networks belonging to different operators. Furthermore, BGP employs mechanisms for loop avoidance and control over the number and type of routes advertised and accepted.

Figure 4 shows an example scenario whereby two BGP peerings are established (for resiliency) between each of the four CsC-CEs (which are actually PE routers of the BGP/MPLS VPN customer carrier network) and a pair of target CsC-PE routers (which are the PE routers of the NGN backbone carrier network).

Figure 4: BGP Plus Labels as the Routing Protocol Between CsC-CEs and CsC-PEs




Label Switching of Customer Packets

As shown in Figure 5, viewing packet flow from left to right, a unicast packet originates as a native IP packet when presented from the end client CE router to the MPLS VPN PE router, which is behaving as a CsC-CE in this context. Upon traversal between CsC-CEs in different MPLS VPN POP locations connected by an NGN backbone carrier using CsC, the packet ultimately undergoes three levels of label encapsulation:

- The innermost label corresponds to the *End Customer VRF*. This label is transparent to the NGN backbone carrier (that is, it is not operated upon in lookup and forwarding tables with the NGN). It is label “A” in Figure 5.
- The middle label is the “outer label” as far as the CsC-CE is concerned, swapped at the CsC-PE, and becomes the “inner label” as far as the NGN backbone carrier is concerned. In Figure 5, this label is assigned as label “B” by the CsC-CE as instructed by the CsC-PE through the BGP plus labels (RFC 3107-compliant) peering. At the CsC-PE itself, the label is swapped (to become label “C” in Figure 5) and is used to associate the packet with the CsC VRF. The packet is then identifiable at the destination CsC-PE at the far end of the backbone carrier network; it allows forwarding to the correct interface.
- The outermost label (shown as label “D” in Figure 5) is assigned by the backbone carrier LDP process at the CsC-PE router, and is present only to allow transport across the backbone carrier CsC core. Thus when a packet leaves the CsC-PE for transport across the backbone carrier core it has three levels of labels on each packet.

Figure 5: Label Encapsulation and End-to-End Packet Flow Across a CsC Core Network



As shown in Figure 5, the last P router in the backbone carrier path has “popped” the outermost label (label “D”) using penultimate-hop label forwarding. The destination CsC-PE uses and removes the middle label (label “C”) to indicate the correct outgoing interface, leaving only the innermost label on presentation to the CsC-CE (label “A”). This CsC-CE, which is the PE router in relation to the end VPN services, uses the last remaining label to determine the VRF route and interface on which to send the native IP packet so that it reaches the required client CE router.

Design and Operational Challenges

The previous section outlined the architectural framework of using CsC to integrate one BGP/MPLS core network with another. This section addresses the important design and operational challenges that such a network transformation brings about.

Maintaining Performance Levels

Many existing operators of “carrier-class” BGP/MPLS networks exploit IP QoS mechanisms to allow different IP-based traffic types to be treated in different ways in terms of how the packets are conveyed across the core network. This treatment relates chiefly to prioritisation of delay, jitter, and/or loss-sensitive traffic, against traffic types that are less sensitive to loss or delay. Customers of VPN services supported on such networks generally demand support of a range of traffic types, including corporate intranet, transactional applications, mail services, data backup, video, and VoIP telephony.

To deal with the range of traffic types, BGP/MPLS VPN service providers have developed the means of supporting IP QoS defining different transport classes with associated service levels. One such example may map, for instance, six service classes based on IETF “Per-Hop Behaviours” as defined by the *Differentiated Services* (DiffServ) working group^[8, 9] and the recommended *DiffServ Code Point* (DSCP) values for them. The classes in this example could be broadly described as follows:

- *Expedited Forwarding* (EF), designed and optimised for the delivery of jitter and delay-sensitive applications such as VoIP
- *Assured Forwarding* (AF), intended to support priority data applications; the AF class is split into four equivalent sub-classes (AF1–AF4) used to segregate data or video traffic applications, with priority being maintained over the Default class
- *Default* (DE), to support “best-effort” (that is, unprioritized) data traffic

The DSCP markings dictate the way in which such traffic is placed into queues and conveyed across the core network. At the edge of the MPLS core, the PE maps the incoming DSCP value into the MPLS *Class-of-Service* (CoS) bits (formerly known as EXP bits).

The details of the mapping relate to the specific implementation and policy of the service provider. Under heavy traffic load and congestion situations, such policies dictate how packets are treated in terms of scheduling, queuing, and discard eligibility.

Both the existing BGP/MPLS “customer carrier” and the target NGN “backbone carrier” networks already have their own implementation of QoS classes to allow management and prioritisation of multiple traffic types carried across their respective core infrastructures. A significant design challenge that arises with integrating the networks is that a suitable mapping of the QoS schema present on the PE routers of the customer carrier network (the CsC-CEs in earlier diagrams) to the QoS schema supported on the PE routers of the NGN (the CsC-PEs in earlier diagrams) is necessary.

It is imperative that such a mapping not compromise the existing customer experience for VPN services in terms of packet loss, packet delay, and packet jitter (that is, delay variance). Careful design, mapping of the required service levels, and ultimately end-to-end testing of the QoS mappings is therefore necessary to assure the maintenance of performance levels after the networks are integrated with CsC.

Network Resiliency

As described earlier in the article and shown in Figure 2, an existing standalone BGP/MPLS network platform has interconnected POP locations using underlying core transmission infrastructures such as SONET/SDH/*Dense Wavelength-Division Multiplexing* (DWDM). The actual number of WAN circuits deployed, the use of transmission-layer protection mechanisms, and the overall topological connectivity between POPs determine overall levels of network resiliency. In turn, this aspect of the network architecture significantly affects the overall level of service availability to end customers of VPN services.

When the standalone BGP/MPLS network has its existing core topology replaced with that of the NGN backbone carrier, it is very important to consider the levels of resiliency delivered with the new integrated core architecture, compared with the existing standalone arrangement. Critical considerations include:

- The physical connectivity between the serving nodes of the customer carrier and the backbone carrier should avoid single points of failure where possible.
- If the physical connectivity between the customer carrier and backbone carrier requires the use of WAN transmission links because locations are geographically separate, then suitable levels of circuit protection should be employed
- Because the backbone carrier effectively replaces the existing core topology of the customer carrier, the actual way in which backbone carrier nodes are interconnected and levels of WAN transmission protection etc., should be analysed.

All these aspects should be assessed and incorporated into the actual design process such that there is no detrimental effect on overall levels of service availability to the end customer. Service levels can be verified by reliability modeling of the new network topology, and by comparing the results with the reliability data for the existing topology.

Fault Management

There are many facets of monitoring and managing a core BGP/MPLS network in terms of assurance of service, alarm detection and filtering, customer notification of faults, and so on. In a standalone network environment, it is generally the responsibility of a particular operational team to manage faults on the network and provide service continuity during various types of failure scenarios. As shown in Figure 6, this operational function usually covers all core network elements, including PE and P (core) routers, as well as the WAN topology interconnecting the service nodes or “POPs.”

In an integrated core network scenario, however, part of the customer carrier network—the P (core) routers and WAN transmission links, for example—are replaced by the NGN backbone carrier. The NGN backbone carrier has its own operational team with specific processes and systems for carrying out monitoring and management of fault events. A crucial challenge arises in terms of how to realise end-to-end fault management holistically and transparently between customer carrier and backbone carrier networks (Figure 6). Important considerations include:

- The requirement for a clear and unambiguous demarcation between customer carrier and backbone carrier core platforms must be addressed in terms of operational responsibility for specific faults and the hand-over procedures between operational domains.
- The use of existing monitoring tools and systems in both the customer carrier and backbone carrier domains must be assessed to determine whether new interfaces between such systems need to be developed to facilitate the hand-over procedures.

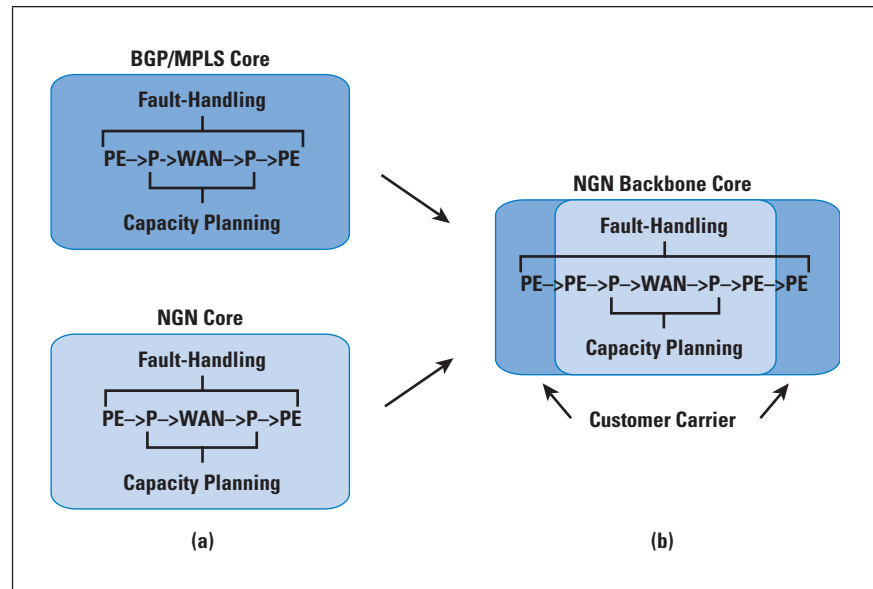
These topics must be factored in to determine the optimal solution for realising smooth and transparent fault-management procedures in an integrated core BGP/MPLS network environment.

Capacity Planning

As shown in Figure 6, in a standalone BGP/MPLS VPN network environment, a particular operational function exists for ongoing core capacity planning to ensure P router and WAN link capacity are suitably dimensioned to cope with current and future traffic demands. When an existing BGP/MPLS VPN network becomes a customer carrier network that is integrated with a target NGN backbone using CsC, there will be a corresponding shift in responsibility for certain aspects of core capacity planning.

VPN service traffic that would have been confined to its own dedicated core network will now be offered onto the NGN backbone carrier core network. As such, the capacity-management function for the NGN backbone carrier must use traffic planning information pertaining to the VPN services in addition to all the other service types supported on the NGN. This aggregated view of traffic demands will accelerate the core capacity dimensioning on the NGN backbone carrier network.

Figure 6: Fault-Management and Capacity-Planning Functions
(a) Before Core Integration
(b) After Core Integration with CsC



Conclusions

The MPLS-based Carrier-supporting-Carrier (CsC) framework provides network operators with a potential solution for integrating an existing BGP/MPLS VPN network, with a target all-IP based NGN. This solution should enable both capital and operational cost reduction by collapsing multiple core networks into a single NGN core domain. The article emphasised that as well as understanding the critical network architectural building blocks required to implement CsC, there are numerous critical design and operational challenges that an integrated core network presents. These challenges include how to maintain service levels and performance metrics for existing VPN customers, resiliency, fault management, and capacity planning. It is important to note, however, that in addition to the broad topic areas covered in this article, many specific additional challenges will present themselves to network operators who have implemented BGP/MPLS VPN networks, and/or NGN networks in their own specific way.

References

- [1] P. Knight and C. Lewis, "Layer 2 and 3 Virtual Private Networks: Taxonomy, Technology, and Standardization Efforts," *IEEE Communications Magazine*, June 2004, pp. 124–131.
- [2] E. Rosen and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)," RFC 4364, February 2006.
- [3] M. Mahmoud, "Carrier-supporting-Carrier: The Whole Story (1)," *Networkers Online*, December 2008.
<http://networkers-online.com/blog/2008/12/carrier-supporting-carrier-the-whole-story-1/>
- [4] M. Mahmoud, "Carrier-supporting-Carrier: The Whole Story (2)," *Networkers Online*, December 2008.
<http://networkers-online.com/blog/2008/12/carrier-supporting-carrier-the-whole-story-2/>
- [5] J. Moy, "OSPF Version 2," RFC 2328, April 1998.
- [6] L. Andersson et al., "LDP Specification," RFC 5036, October 2007.
- [7] Y. Rekhter and E. Rosen, "Carrying Label Information in BGP-4," RFC 3107, January 2001.
- [8] B. Davy et al., "An Expedited Forwarding PHB," RFC 3246, March 2002.
- [9] J. Heinanen et al., "Assured Forwarding PHB Group," RFC 2597, June 1999.

PAUL VEITCH holds an M.Eng. and a Ph.D. from the University of Strathclyde, Glasgow. He joined BT at Martlesham Heath, Ipswich, UK, in September 1996, and worked on various aspects of broadband transmission architectures, multi-service platforms, and 3G network design. In 2000, he joined MCI-WorldCom in Cambridge, UK, and led a number of projects on IP backbone network design. In 2003, he returned to BT to work on IP VPN infrastructure design. He is currently the design authority for BT Retail's Internet networks. He can be reached at: paul.veitch@bt.com

PAUL HITCHEN holds a B.Eng. in Electrical and Electronic Engineering from the University of Salford. He joined BT at Martlesham Heath in September 1990 and has worked on numerous aspects of BT's data services. From 1990 to 1997 he led the development of BT's multiprotocol router portfolio, developing routing and QoS functions with BT's equipment suppliers and provided consulting to BT's customers on IP and Ethernet networks. During the same period he worked on the introduction of Frame Relay, ATM, and SMDS WAN services for BT. In 1997 he developed BT's first IP VPN service offering, working on the development and standardisation of MPLS and VPN technology. From 1997 to the end of 2006 he led the design of BT's Global MPLS Network and service, expanding the network to provide service to more than 150 countries across the world. He is currently a principal consultant working on BT's 21CN IP/MPLS network, focusing on the integration of BT's networks onto 21CN and introducing content delivery and IPTV into the network. He can be reached at: paul.hitchen@bt.com

MARTIN MITCHELL holds an M.Sci. from the University of Bristol and has worked for BT since 2007. He is currently an IP network designer specializing in service provider core design, Ethernet access, and network migrations. He can be reached at: martin.3.mitchell@bt.com

Letter to the Editor

Hi Ole,

I enjoyed the article entitled “PMIPv6: A Network-Based Localized Mobility Management Solution” in the last issue of *The Internet Protocol Journal* (Volume 13, No. 3, September 2010).

I believe that in the “Security Considerations” section it should be mentioned that the CSI (Cga & Send maintenance) working group in the IETF is also working on updating the *Secure Neighbor Discovery* (SEND) specification (RFC 3971) to include the possibility of authenticating the proxied *Neighbor Discovery* (ND) messages sent between the terminal, the *Mobile Access Gateway* (MAG), and the *Local Mobility Anchor* (LMA). This configuration should work in addition to the proposed *IP Security* (IPsec) tunnel between the MAG and the LMA.

The reference material is available at:

<https://datatracker.ietf.org/doc/draft-ietf-csi-proxy-send/>

<https://datatracker.ietf.org/doc/draft-ietf-csi-send-cert/>

Regards,

—Roque Gagliano, Cisco Systems
rogaglia@cisco.com

One of the authors responds:

Dear Ole and Roque,

Thanks for reading our article and providing these valuable comments. We agree with your point. We just considered the basic security mechanisms in our article, limiting the scope to the protocols already standardized, which cover only the protection of the MAG-LMA signaling. We agree that the efforts being carried out within the CSI working group are worth mentioning with regard to the security aspects of PMIPv6.

Thanks,

—Carlos J. Bernardos, Universidad Carlos III de Madrid
cjb@it.uc3m.es

Book Review

A History of the Internet

A History of the Internet and the Digital Future, by Johnny Ryan, Reaktion Books, ISBN 978 1 86189 777 0, September 2010.

Any attempt to document a 50-year history of people and activities that had such a profound and global effect as the Internet faces some challenges. Sequences are complex; written source materials are sketchy; and the many different memories conflict. Added to this reality, of course, are legitimate disagreements about intents and effects. To evaluate such writing effort means first looking for useful criteria. Here are mine: In terms of basic research, was the effort extensive, looking for multiple, appropriate sources and exploring a wide range of probing and constructive questions? Were the sources and questions interesting? This line of thinking leads to a query about the way the author integrates the resulting massive body of data. Is there an effort to develop critical analyses? Are alternative explanations explored?

Johnny Ryan's ambitious *A History of the Internet and the Digital Future* is a rather modest 246 pages, including 28 pages of references. Overall my feeling is that he does quite an interesting job of satisfying the first half of his title, but a somewhat disappointing job with the second half. His research was extensive throughout, but he takes a more critical view of the history than he does of the social aspects of our digital future. In the first half, he integrates information and reports discrepancies and curiosities. In the second half, he indulges in the common, wide-eyed wonderment that technology futurist efforts inherently risk. (Full disclosure: By way of demonstrating the thoroughness of his research, Ryan even included me as one of his many sources.)

Organization

The book is divided into three parts. Broadly, they cover origins, growth, and social effects. Ryan's use of "centrifugal" is contrasted with "centripetal" and is meant to distinguish paradigmatic tensions between approaches that centralize control versus approaches that distribute it. (Oddly, neither of these pivotal terms is in the index.) On page 8 he sets the stage:

"Three characteristics have asserted themselves throughout the Internet's history and will define the digital age to which we must all adjust: The Internet is a centrifugal force, user-driven and open."

By "centrifugal" he means moving outward, away from centralized control. For me, the terminology proved distracting, because I kept hearing my 8th-grade science teacher condescendingly explaining that there is no physics force called centrifugal. Rather it is a perception of the interaction between inertia and centripetal force.

For those with less compulsive (or effective) science teachers, the analogy might prove more helpful, because the design choice really is central to the history of networking. The tension between centralized versus distributed has marked—and continues to mark—much of the development of networking. In fact, I wish Ryan had explored its continuation as much as he explored its effect on origins.

Early History

In general, Ryan presents a narrative with fine-grained detail of the different players who played a critical role in the creation and pursuit of packet switching and then its evolution to link independent networks and technologies^[1]. Efforts to take credit for the former have often become quite public and unseemly; Ryan dissects the play of actors, the essence of their technical ideas, and the details of their activities with documentation and diligence, and even uncovers some discrepancies. He develops a narrative that I found intriguing, enlightening, and credible. What I especially liked was that he explored the organizational milieu in which the activities took place. So we hear of the origins of groups such as the *Advanced Research Projects Agency* (ARPA), Lincoln Labs, and The Rand Corporation; the social and political forces that created them; and the roles they played.

Narrative Arcs

The following is really the strength of this book: It develops narrative arcs about social, political, and organizational environments and the steps taken within them that moved along the path of the Internet. It explores who, when, how, and what, both overall and in detail. At its best, the book provides comparative perspective to help the reader understand what was risky and truly innovative and thereby understand what was really challenging to develop and get adopted. As a minor example, Ryan deserves credit for his exploration and debunking of the media distortions surrounding Al Gore's role and statements concerning the Internet. Strictly speaking, debunking media excesses would not normally seem relevant to a review of the history of a technology, but Ryan uses this example for some consideration of the role of politics in the development of the Internet. The U.S. government could have chosen to assume more control over the Internet; it might have quickly turned it into a telecommunications monopoly, rather than letting it develop through independent market forces.

As would be expected for a story this sweeping, Ryan is sometimes redundant and sometimes inconsistent. Overall, the book would have benefited from more careful editing. So it has a quick reference to the “invention” of e-mail messaging at Bolt Beranek and Newman, but later has a more accurate, detailed account of Ray Tomlinson's 1971 effort, there, to add networking to the *existing* e-mail mechanism. (E-mail messaging was present on the first time-sharing systems of the 1960s, but these systems were standalone services. Tomlinson got them to talk each other.)

Another touchstone I use for discussions of Internet history is the role of the *Computer Science Network* (CSNet), because I worked on that. CSNet served as the forerunner of the larger and more obviously pivotal *National Science Foundation Network* (NSFNet). With NSFNet the Internet developed the ability to support multiple backbones—essential for a truly competitive Internet—and the market-priming creation of regional operational services, from which the seeds of the commercial Internet were sown. Ryan notes the role of CSNet as a kind of market research that led to NSFnet, and in this observation his discussion is notable. But his account of CSNet details is somewhat skewed, because CSNet is cast as having full packet-level connectivity, with e-mail-only telephone-based linkages as a secondary service. In reality full connectivity came later; the original years of CSNet were e-mail-only. Why this fact is important to note—besides overly personal fault-finding—is as a reminder that the accounting efforts for this sort of history are always noisy; the story signal is never pure, even with a diligent effort.

A further touchstone topic is the *Domain Name System* (DNS) and the development of the *Internet Corporation for Assigned Names and Numbers* (ICANN). The interesting part of this saga is later-stage Internet history, and Ryan is relatively sloppy with the details. For example, he muddles what *generic Top-Level Domains* (gTLD) already existed and what new ones were proposed, such as **.com** versus **.biz**; he also muddles the distinction between gTLDs and national domains, such as **.uk**. On the other hand, he certainly captures the continuing tone of controversy that surrounded the development and operation of ICANN, the organization now managing assignment of IP addresses and domain names.

But the most obvious, later-stage touchstone for a history like this one must be the development of the World Wide Web. Ryan gets mixed marks here. He misses the long history of open document publishing that existed even in the earlier *Advanced Research Projects Agency Network* (ARPANET), with “anonymous” FTP, and he misses that the use of *Gopher* predated the web by several years. He also misses just how complete and useful a “dynamically linked document” system Doug Englebart’s NLS (computer) system provided 20 years before the invention of the web^[2]. Hence, he misses the long, historical arc for publishing on the Internet. On the other hand, he does discuss *Gopher* and explores some of the reasons it lost the competition to the web. He focuses on management and intellectual property issues, whereas I tend to consider *Gopher* as having a much poorer cost/benefit mix. *Gopher* was text-only and required going down a potential long lookup tree—quite a few “clicks”—before getting any content. The web is mixed-media and can provide utility to the reader—that is, content—at each step down a lookup path. So the web is more complex to develop than *Gopher*, but it provides enough additional power and better human factors to be worth it.

Ryan's discussion of the commercial explosive growth of the Internet is a good read, including the Dutch tulip market reference and his introduction to some relevant tidbits of economics theory. However, as the book moves into "Web 2.0" and beyond, it provides reasonable descriptions of who did what to create popular new services, but his critical eye largely stops providing serious analysis. Explanations sound more like exuberance than examination. On the other hand, he certainly provides substance to the view that the Internet enables "long-tail" market opportunities to discover and satisfy specialized segments. His discussion of politicians' inventive use of the Internet is nicely concise and integrated. Again, it provides a narrative arc with substance. But his predictions for the future of users as news consumers or as citizens in political processes have too much tone of certitude and positive outcome than is justifiable in my opinion.

Worth Reading

In sum, the book is certainly worth reading. You will likely learn quite a bit, but make sure you read with glasses that have no hint of rose coloring!

References

- [1] Debating which milestone marks "the beginning of the Internet" is a favorite pastime, including among those around during the period in question. Various definitions are legitimate, as long as one is clear about the choice. For me, the operational demonstration of packet switching was when the world changed, so I choose 1969 and the first four nodes of the ARPANET; or its public demonstration in 1972. TCP/IP built on this, by refining and minimizing the work to be done within the infrastructure and by linking independent networks.
- [2] In the early 1970s, my job at UCLA included technical documentation and supporting online use by the Computer Science Department's secretaries. We did all our editing remotely, on the Engelbart system, because it was so powerful.

—Dave Crocker, *Brandenburg Internet Working*
dcrocker@bbiw.net

Read Any Good Books Lately?

Then why not share your thoughts with the readers of IPJ? We accept reviews of new titles, as well as some of the "networking classics." In some cases, we may be able to get a publisher to send you a book for review if you don't have access to it. Contact us at ipj@cisco.com for more information.



Photo: Matsuzaki Yoshinobu

Bjoern A. Zeeb Receives Second Itojun Service Award

The second Itojun Service Award was presented at the 79th meeting of the *Internet Engineering Task Force* (IETF) in Beijing, China. Bjoern A. Zeeb received the award for his dedicated work to make significant improvements in open source implementations of IPv6.

First awarded last year, the *Itojun Service Award* honours the memory of Dr. Jun-ichiro “Itojun” Hagino, who passed away in 2007, aged just 37. The award, established by the friends of Itojun and administered by the *Internet Society* (ISOC), recognises and commemorates the extraordinary dedication exercised by itojun over the course of IPv6 development.

“For many years, Bjoern has been a committed champion of, and contributor to, implementing IPv6 in open source operating systems used in servers, desktops, and embedded computer platforms, including those used by some of the busiest websites in the world,” said Jun Murai of the Itojun Service Award Committee and Founder of the WIDE Project. “On behalf of the Itojun Service Award Committee, I am extremely pleased to present this award to Bjoern for his outstanding work in support of IPv6 development and deployment.”

The Itojun Service Award is focused on pragmatic contributions to developing and deploying IPv6 in the spirit of serving the Internet. The award, expected to be presented annually, includes a presentation crystal, a US\$3,000 honorarium, and a travel grant.

“This is a great honour, and I would like to thank the people who recommended me for the award and the committee for believing my work was valuable. I never met Itojun but he was one of the people helping me, and I have the highest respect for his massive foundational work,” said Bjoern A. Zeeb. “As the Internet community works to roll out IPv6 to more and more people all around the globe, we also need to help others—developers, businesses, and users—understand and use the new Internet protocols so that the vision Itojun was working so hard for comes true.”

Each Internet-connected device uses an IP address and, with the number of Internet-connected devices growing rapidly, the supply of unallocated IPv4 addresses is expected to be exhausted within the next year. To help ensure the continued rapid growth of the Internet, IPv6 provides a huge increase in the number of available addresses. And, while the technical foundations of IPv6 are well established, significant work remains to expand the deployment and use of IPv6.

For more information about the Itojun Service Award see:
<http://www.isoc.org/itojun/>

Remaining IPv4 Address Space Drops Below 5 percent

The *Number Resource Organization* (NRO) recently announced that less than five percent of the world's IPv4 addresses remain unallocated. APNIC, the Regional Internet Registry for the Asia Pacific region, has been assigned two blocks of IPv4 addresses by the *Internet Assigned Numbers Authority* (IANA). This latest allocation means that the IPv4 free pool dipped below 10% in January 2010. Since then, over 200 million IPv4 addresses have been allocated from IANA to the *Regional Internet Registries* (RIRs).

“This is a major milestone in the life of the Internet, and means that allocation of the last blocks of IPv4 to the RIRs is imminent,” stated Axel Pawlik, Chairman of the NRO, the official representative of the five RIRs. “It is critical that all Internet stakeholders take definitive action now to ensure the timely adoption of IPv6.”

IPv6 is the “next generation” of the Internet Protocol, providing a hugely expanded address space, which will allow the Internet to grow into the future. In 2010, the five RIRs are expected to allocate over 2,000 IPv6 address blocks, representing an increase of over 70% on the number of IPv6 allocations in 2009. In contrast, the number of IPv4 allocations is expected to grow by only 8% in 2010. These statistics indicate an absence of any last minute “rush” on IPv4 addresses, and a strong momentum behind the adoption of IPv6.

“The allocation of Internet number resources by the five RIRs enables every region in the world to benefit from fair and equitable distribution of IPv4 and IPv6 addresses. We are also actively collaborating with stakeholders at the local, regional, and global level to offer training and advice to public and private sector organisations on IPv6 adoption to ensure that everyone is prepared for IPv4 depletion and IPv6 deployment,” added Pawlik.

The IANA assigns IPv4 addresses to the RIRs in blocks that equate to 1/256th of the entire IPv4 address space (each block is referred to as a “/8” or “slash-8”). The most recent assignment means that there are now only 12 of these blocks available, which is less than five percent of the entire IPv4 address pool.

The final five blocks of IPv4 addresses will be distributed simultaneously to the five RIRs, leaving only seven blocks to be handed out under the normal distribution method.

According to current depletion rates, the last five IPv4 address blocks will be allocated to the RIRs in early 2011. The pressure to adopt IPv6 is mounting. Many worry that without adequate preparation and action, there will be a chaotic scramble for IPv6, which could increase Internet costs and threaten the stability and security of the global network.

The NRO exists to protect the pool of unallocated Internet numbers (IP addresses and AS numbers) and serves as a coordinating mechanism for the five RIRs to act collectively on matters relating to the interests of RIRs. For further information, visit <http://www.nro.net>

The RIRs are independent, not-for-profit membership organizations that support the infrastructure of the Internet through technical coordination. There are five RIRs in the world today. Currently, the IANA allocates blocks of IP addresses and ASNs, known collectively as *Internet Number Resources*, to the RIRs, who then distribute them to their members within their own specific service regions. RIR members include *Internet Service Providers* (ISPs), telecommunications organizations, large corporations, governments, academic institutions, and industry stakeholders, including end users.

The RIR model of open, transparent participation has proven successful at responding to the rapidly changing Internet environment. Each RIR holds one to two open meetings per year, as well as facilitating online discussion by the community, to allow the open exchange of ideas from the technical community, the business sector, civil society, and government regulators. Each RIR performs a range of critical functions including: The reliable and stable allocation of Internet number resources (IPv4, IPv6 and *Autonomous System Number* resources); The responsible storage and maintenance of this registration data; The provision of an open, publicly accessible database where this data can be accessed. RIRs also provide a range of technical and coordination services for the Internet community. The five RIRs are:

AfriNIC: <http://www.afrinic.net>

APNIC: <http://www.apnic.net>

ARIN: <http://www.arin.net>

LACNIC: <http://www.lacnic.net>

RIPE NCC: <http://www.ripe.net>

The Facebook logo, consisting of the word "facebook" in white lowercase letters on a blue rectangular background.

Find us on Facebook

In addition to *The Internet Protocol Forum*, available at <http://www.ipjforum.org>, IPJ now has its own Facebook page. Join the discussion and get the latest news and updates:

<http://www.facebook.com/#!/pages/Internet-Protocol-Journal/163288673690055>

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
Tel: +1 408 526-4000
E-mail: ipj@cisco.com*

*Copyright © 2010 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.

