Network Working Group
Request for Comments: 4958
Category: Informational

K. Carlberg
G11

July 2007

A Framework for Supporting Emergency Telecommunications Services (ETS) within a Single Administrative Domain

Status of This Memo

Copyright Notice

Abstract

This document presents a framework discussing the role of various protocols and mechanisms that could be considered candidates for supporting Emergency Telecommunication Services (ETS) within a single administrative domain. Comments about their potential usage as well as their current deployment are provided to the reader. Specific solutions are not presented.

Table of Contents

1. Introduction

This document presents a framework for supporting Emergency Telecommunications Services (ETS) within the scope of a single administrative domain. This narrow scope provides a reference point for considering protocols that could be deployed to support ETS. [rfc4375] is a complementary effort that articulates requirements for a single administrative domain and defines it as "collection of resources under the control of a single administrative authority". We use this other effort as both a starting point and guide for this document.

A different example of a framework document for ETS is [rfc4190], which focused on support for ETS within IP telephony. In this case, the focal point was a particular application whose flows could span multiple autonomous domains. Even though this document uses a somewhat more constrained perspective than [rfc4190], we can still expect some measure of overlap in the areas that are discussed.

1.1. Differences between Single and Inter-Domain

The progression of our work in the following sections is helped by stating some key differences between the single and inter-domain cases. From a general perspective, one can start by observing the following.

a) Congruent with physical topology of resources, each domain is an authority zone, and there is currently no scalable way to transfer authority between zones.

b) Each authority zone is under separate management.

c) Authority zones are run by competitors; this acts as further deterrent to transferring authority.

As a result of the initial statements in (a) through (c) above, additional observations can be made that distinguish the single and inter-domain cases, as follows.

d) Different policies might be implemented in different administrative domains.

e) There is an absence of any practical method for ingress nodes of a transit domain to authenticate all of the IP network layer packets that have labels indicating a preference or importance.

f) Given item (d) above, all current inter-domain QoS mechanisms at the network level generally create easily exploited and significantly painful Denial of Service (DoS) / Distributed Denial of Service (DDoS) attack vectors on the network.

g) A single administrative domain can deploy various mechanisms (e.g., access control lists) into each and every edge device (e.g., ethernet switch or router) to ensure that only authorized end-users (or layer 2 interfaces) are able to emit frames/packets with non-default QoS labels into the network. This is not feasible in the inter-domain case because the inter-domain link contains aggregated flows. In addition, the dissemination of access control lists at the network level is not scalable in the inter-domain case.

h) A single domain can deploy mechanisms into the edge devices to enforce its domain-wide policies -- without having to trust any third party to configure things correctly. This is not possible in the inter-domain case.

While the above is not an all-inclusive set of differences, it does provide some rationale why one may wish to focus efforts in the more constrained scenario of a single administrative domain.

2. Common Practice: Provisioning

The IEPREP working group and mailing list have had extensive discussions about over-provisioning. Many of these exchanges have debated the need for QoS mechanisms versus over-provisioning of links.

In reality, most IP network links are provisioned with a percentage of excess capacity beyond that of the average load. The ’shared’ resource model together with TCP’s congestion avoidance algorithms helps compensate for those cases where spikes or bursts of traffic are experienced by the network.

The thrust of the debate within the IEPREP working group is whether it is always better to over-provision links to such a degree that spikes in load can still be supported with no loss due to congestion. Advocates of this position point to many ISPs in the US that take this approach instead of using QoS mechanisms to honor agreements with their peers or customers. These advocates point to cost effectiveness in comparison to complexity and security issues associated with other approaches.

Proponents of QoS mechanisms argue that the relatively low cost of bandwidth enjoyed in the US (particularly, by large ISPs) is not necessarily available throughout the world. Beyond the subject of cost, some domains are comprised of physical networks that support wide disparity in bandwidth capacity -- e.g., attachment points connected to high capacity fiber and lower capacity wireless links.

This document does not advocate one of these positions over the other. The author does advocate that network administrators/operators should perform a cost analysis between over-provisioning for spikes versus QoS mechanisms as applied within a domain and its access link to another domain (e.g., a customer and its ISP). This analysis, in addition to examining policies and requirements of the administrative domain, should be the key to deciding how (or if) ETS should be supported within the domain.

If the decision is to rely on over-provisioning, then some of the following sections will have little to no bearing on how ETS is supported within a domain. The exception would be labeling mechanisms used to convey information to other communication architectures (e.g., SIP-to-SS7/ISUP gateways).

3. Objective

The primary objective is to provide a target measure of service within a domain for flows that have been labeled for ETS. This level may be better than best effort, the best available service that the network (or parts thereof) can offer, or a specific percentage of resource set aside for ETS. [rfc4375] presents a set of requirements in trying to achieve this objective.

This framework document uses [rfc4375] as a reference point in discussing existing areas of engineering work or protocols that can play a role in supporting ETS within a domain. Discussion of these areas and protocols are not to be confused with expectations that they exist within a given domain. Rather, the subjects discussed in Section 4 below are ones that are recognized as candidates that can exist and could be used to facilitate ETS users or data flows.

3.1. Scenarios

One of the topics of discussion on the IEPREP mailing list and in the working group meetings is the operating environment of the ETS user. Many variations can be dreamed of with respect to underlying network technologies and applications. Instead of getting lost in hundreds of potential scenarios, we attempt to abstract the scenarios into two simple case examples.

(a) A user in their home network attempts to use or leverage any ETS capability within the domain.

(b) A user visits a foreign network and attempts to use or leverage any ETS capability within the domain.

We borrow the terms "home" and "foreign" network from that used in Mobile IP [rfc3344]. Case (a) is considered the normal and vastly most prevalent scenario in today's Internet. Case (b) above may simply be supported by the Dynamic Host Configuration Protocol (DHCP) [rfc2131], or a static set of addresses, that are assigned to 'visitors' of the network. This effort is predominantly operational in nature and heavily reliant on the management and security policies of that network.

A more ambitious way of supporting the mobile user is through the use of the Mobile IP (MIP) protocol. MIP offers a measure of application-transparent mobility as a mobile host moves from one subnetwork to another while keeping the same stable IP address registered at a global anchor point. However, this feature may not always be available or in use. In any case, where it is in use, at least some of the packets destined to and from the mobile host go through the home network.

4. Topic Areas

The topic areas presented below are not presented in any particular order or along any specific layering model. They represent capabilities that may be found within an administrative domain. Many are topics of on-going work within the IETF.

It must be stressed that readers of this document should not expect any of the following to exist within a domain for ETS users. In many cases, while some of the following areas have been standardized and in wide use for several years, others have seen very limited deployment.

4.1. MPLS

Multiprotocol Label Switching (MPLS) is generally the first protocol that comes to mind when the subject of traffic engineering is brought up. MPLS signaling produces Labeled Switched Paths (LSPs) through a network of Label Switch Routers [rfc3031]. When traffic reaches the ingress boundary of an MPLS domain (which may or may not be congruent with an administrative domain), the packets are classified, labeled, scheduled, and forwarded along an LSP.

[rfc3270] describes how MPLS can be used to support Differentiated Services. The RFC discusses the use of the 3-bit EXP (experimental) field to convey the Per Hop Behavior (PHB) to be applied to the packet. As we shall see in later sections, this 3-bit field can be mapped to fields in several other protocols.

The inherent features of classification, scheduling, and labeling are viewed as symbiotic, and therefore, they are often integrated with other protocols and architectures. Examples of this include RSVP and Differentiated Services. Below, we discuss several instances where a given protocol specification or mechanism has been known to be complemented with MPLS. This includes the potential labels that may be associated with ETS. However, we stress that MPLS is only one of several approaches to support traffic engineering. In addition, the complexity of the MPLS protocol and architecture may make it suited only for large domains.

4.2. RSVP

The original design of RSVP, together with the Integrated Services model, was one of an end-to-end signaling capability to set up a path of reserved resources that spanned networks and administrative domains [rfc2205]. Currently, RSVP has not been widely deployed by network administrators for QoS across domains. Today’s limited deployment by network administrators has been mostly constrained to boundaries within a domain, and commonly in conjunction with MPLS signaling. Early deployments of RSVP ran into unanticipated scaling issues; it is not entirely clear how scalable an RSVP approach would be across the Internet.

[rfc3209] is one example of how RSVP has evolved to complement efforts that are scoped to operate within a domain. In this case, extensions to RSVP are defined that allow it to establish intra-domain Labeled Switched Paths (LSPs) in Multiprotocol Label Switching (MPLS).

[rfc2750] specifies extensions to RSVP so that it can support generic policy-based admission control. This standard goes beyond the support of the POLICY_DATA object stipulated in [rfc3209], by defining the means of control and enforcement of access and usage policies. While the standard does not advocate a particular policy architecture, the IETF has defined one that can complement [rfc2750] -- we expand on this in Section 4.3 below.

4.2.1. Relation to ETS

The ability to reserve resources correlates to an ability to provide preferential service for specifically classified traffic -- the classification being a tuple of 1 or more fields which may or may not include an ETS specific label. In cases where a tuple includes a label that has been defined for ETS usage, the reservation helps ensure that an emergency-related flow will be forwarded towards its destination. Within the scope of this document, this means that RSVP would be used to facilitate the forwarding of traffic within a domain.

We note that this places an importance on defining a label and an associated field that can be set and/or examined by RSVP-capable nodes.

Another important observation is that major vendor routers currently constrain their examination of fields for classification to the network and transport layers. This means that application layer labels will mostly likely be ignored by routers/switches.

4.3. Policy

The Common Open Policy Service (COPS) protocol [rfc2748] was defined to provide policy control over QoS signaling protocols, such as RSVP. COPS is based on a query/response model in which Policy Enforcement Points (PEPs) interact with Policy Decision Points (i.e., policy servers) to exchange policy information. COPS provides application-level security and can operate over IPsec or TLS. COPS is also a stateful protocol that supports a push model. This means that servers can download new policies or alter existing ones to known clients.

[rfc2749] articulates the usage of COPS with RSVP. It specifies COPS client types, context objects, and decision objects. Thus, when an RSVP reservation is received by a PEP, the PEP decides whether to accept or reject it based on policy. This policy information can be stored a priori to the reception of the RSVP PATH message, or it can be retrieved on an on-demand basis. A similar course of action could be applied in cases where ETS-labeled control flows are received by the PEP. This of course would require an associated (and new) set of documents that first articulates types of ETS signaling and then specifies its usage with COPS.

A complementary document to the COPS protocols is COPS Usage for Policy Provisioning (COPS-PR) [rfc3084].

As a side note, the current lack of deployment by network administrators of RSVP has also played at least an indirect role in the subsequent lack of implementation and deployment of COPS-PR. [rfc3535] is an output from the IAB Network Management Workshop in which the topic of COPS and its current state of deployment was discussed. At the time of that workshop in 2002, COPS-PR was considered a technology/architecture that did not fully meet the needs of network operators. It should also be noted that at the 60th IETF meeting held in San Diego in 2004, COPS was discussed as a candidate protocol that should be declared as historic because of lack of use and concerns about its design. In the future, an altered design of COPS may emerge that addresses the concern of operators, but speculation on that or other possibilities is beyond the scope of this document.

4.4. Subnetwork Technologies

This is a generalization of work that is considered "under" IP and for the most part outside of the IETF standards body. We discuss some specific topics here because there is a relationship between them and IP in the sense that each physical network interacts at its edge with IP.

4.4.1. IEEE 802.1 VLANs

The IEEE 802.1q standard defined a tag appended to a Media Access Controller (MAC) frame for support of layer 2 Virtual Local Area Networks (VLANs). This tag has two parts: a VLAN identifier (12 bits) and a Prioritization field of 3 bits. A subsequent standard, IEEE 802.1p, later incorporated into a revision of IEEE 802.1d, defined the Prioritization field of this new tag [iso15802]. It consists of 8 levels of priority, with the highest priority being a value of 7. Vendors may choose a queue per priority codepoint, or aggregate several codepoints to a single queue.

The 3-bit Prioritization field can be easily mapped to the old ToS field of the upper-layer IP header. In turn, these bits can also be mapped to a subset of differentiated codepoints. Bits in the IP header that could be used to support ETS (e.g., specific Diffserv codepoints) can in turn be mapped to the Prioritization bits of 802.1p. This mapping could be accomplished in a one-to-one manner between the 802.1p field and the IP ToS bits, or in an aggregate manner if one considers the entire Diffserv field in the IP header. In either case, because of the scarcity of bits, ETS users should expect that their traffic will be combined or aggregated with the same level of priority as some other types of "important" traffic. In other words, given the existing 3-bit Priority Field for 802.1p, there will not be an exclusive bit value reserved for ETS traffic.

Certain vendors are currently providing mappings between the 802.1p field and the ToS bits. This is in addition to integrating the signaling of RSVP with the low-level inband signaling offered in the Priority field of 802.1p.

It is important to note that the 802.1p standard does not specify the correlation of a layer 2 codepoint to a physical network bandwidth reservation. Instead, this standard provides what has been termed as "best effort QoS". The value of the 802.1p Priority codepoints is realized at the edges: either as the MAC payload is passed to upper layers (like IP), or as it is bridged to other physical networks like Frame Relay. Either of these actions help provide an intra-domain wide propagation of a labeled flow for both layer 2 and layer 3 flows.

4.4.2. IEEE 802.11e QoS

The 802.11e standard is a proposed enhancement that specifies mechanisms to provide QoS to the 802.11 family of protocols for wireless LANs.

Previously, 802.11 had two modes of operation. One was Distributed Coordination Function (DCF) , which is based on the classic collision detection schema of "listen before sending". A second optional mode is the Point Coordination Function (PCF). The modes splits access time into contention-free and contention-active periods -- transmitting data during the former.

The 802.11e standard enhances DCF by adding support for 8 different traffic categories or classifications. Each higher category waits a little less time than the next lower one before it sends its data.

In the case of PCF, a Hybrid Coordination Function has been added that polls stations during contention-free time slots and grants them a specific start time and maximum duration for transmission. This second mode is more complex than enhanced DCF, but the QoS can be more finely tuned to offer specific bandwidth and jitter control. It must be noted that neither enhancement offers a guarantee of service.

4.4.3. Cable Networks

The Data Over Cable Service Interface Specification (DOCSIS) is a standard used to facilitate the communication and interaction of the cable subnetwork with upper-layer IP networks [docsis]. Cable subnetworks tend to be asynchronous in terms of data load capacity: typically, 27 M downstream, and anywhere from 320 kb to 10 M upstream (i.e., in the direction of the user towards the Internet).

The evolution of the DOCSIS specification, from 1.0 to 1.1, brought about changes to support a service other than best effort. One of the changes was indirectly added when the 802.1d protocol added the Priority field, which was incorporated within the DOCSIS 1.1 specification. Another change was the ability to perform packet fragmentation of large packets so that Priority-marked packets (i.e., packets marked with non-best effort labels) can be multiplexed in between the fragmented larger packet.

It's important to note that the DOCSIS specifications do not specify how vendors implement classification, policing, and scheduling of traffic. Hence, operators must rely on mechanisms in Cable Modem Termination Systems (CMTS) and edge routers to leverage indirectly or directly the added specifications of DOCSIS 1.1. As in the case of 802.1p, ETS-labeled traffic would most likely be aggregated with other types of traffic, which implies that an exclusive bit (or set of bits) will not be reserved for ETS users. Policies and other managed configurations will determine the form of the service experienced by ETS labeled traffic.

Traffic engineering and management of ETS labeled flows, including its classification and scheduling at the edges of the DOCSIS cloud, could be accomplished in several ways. A simple schema could be based on non-FIFO queuing mechanisms like class-based weighted fair queuing (or combinations and derivations thereof). The addition of active queue management like Random Early Detection could provide simple mechanisms for dealing with bursty traffic contributing to congestion. A more elaborate scheme for traffic engineering would include the use of MPLS. However, the complexity of MPLS should be taken into consideration before its deployment in networks.

4.5. Multicast

Network layer multicast has existed for quite a few years. Efforts such as the Mbone (multicast backbone) have provided a form of tunneled multicast that spans domains, but the routing hierarchy of the Mbone can be considered flat and non-congruent with unicast routing. Efforts like the Multicast Source Discovery Protocol [rfc3618] together with the Protocol Independent Multicast - Sparse Mode (PIM-SM) have been used by a small subset of Internet Service Providers to provide forms of inter-domain multicast [rfc4601]. However, network layer multicast has not been accepted as a common production level service by a vast majority of ISPs.

In contrast, intra-domain multicast in domains has gained more acceptance as an additional network service. Multicast can produce denial-of-service attacks using the any sender model, with the problem made more acute with flood and prune algorithms. Source-

specific multicast [rfc3569], together with access control lists of who is allowed to be a sender, reduces the potential and scope of such attacks.

4.5.1. IP Layer

The value of IP multicast is its efficient use of resources in sending the same datagram to multiple receivers. An extensive discussion on the strengths of and concerns about multicast is outside the scope of this document. However, one can argue that multicast can very naturally complement the push-to-talk feature of land mobile radio (LMR) networks.

Push-to-talk is a form of group communication where every user in the "talk group" can participate in the same conversation. LMR is the type of network used by First Responders (e.g., police, firemen, and medical personnel) that are involved in emergencies. Currently, certain vendors and providers are offering push-to-talk service to the general public in addition to First Responders. Some of these systems are operated over IP networks or are interfaced with IP networks to extend the set of users that can communicate with each other. We can consider at least a subset of these systems as either closed IP networks, or domains, since they do not act as transits to other parts of the Internet.

The potential integration of LMR talk groups with IP multicast is an open issue. LMR talk groups are established in a static manner with man-in-the-loop participation in their establishment and teardown. The seamless integration of these talk groups with multicast group addresses is a feature that has not been discussed in open forums.

4.5.2. IEEE 802.1d MAC Bridges

The IEEE 802.1d standard specifies fields and capabilities for a number of features. In Section 4.3.2 above, we discussed its use for defining a Prioritization field. The 802.1d standard also covers the topic of filtering MAC layer multicast frames.

One of the concerns about multicast is that broadcast storms can arise and generate a denial of service against other users/nodes. Some administrators purposely filter out multicast frames in cases where the subnetwork resource is relatively small (e.g., 802.11 networks). Operational considerations with respect to ETS may wish to consider doing this on an as-needed basis, balancing the conditions of the network with the perceived need for multicast. In cases where filtering out multicast can be activated dynamically, COPS may be a good means of providing consistent domain-wide policy control.

4.6. Discovery

If a service is being offered to explicitly support ETS, then it would seem reasonable that discovery of the service may be of benefit. For example, if a domain has a subset of servers that recognize ETS-labeled traffic, then dynamic discovery of where these servers are (or even if they exist) would be more beneficial than relying on statically configured information.

The Service Location Protocol (SLP) [rfc2608] is designed to provide information about the existence, location, and configuration of networked services. In many cases, the name of the host supporting the desired service is needed to be known a priori in order for users to access it. SLP eliminates this requirement by using a descriptive model that identifies the service. Based on this description, SLP then resolves the network address of the service and returns this information to the requester. An interesting design element of SLP is that it assumes that the protocol is run over a collection of nodes that are under the control of a single administrative authority. This model follows the scope of this framework document. However, the scope of SLP may be better suited for parts of an enterprise network versus an entire domain.

Anycasting [rfc1546] is another means of discovering nodes that support a given service. Interdomain anycast addresses, propagated by BGP, represent anycast in a wide scope and have been used by multiple root servers for a while. Anycast can also be realized in a more constrained and limited scope (i.e., solely within a domain or subnet), and as in the case of multicast, it may not be supported.

[rfc4291] covers the topic of anycast addresses for IPv6. Unlike SLP, users/applications must know the anycast address associated with the target service. In addition, responses to multiple requests to the anycast address may come from different servers. Lack of response (not due to connectivity problems) correlates to the discovery that a target service is not available. Detailed tradeoffs between this approach and SLP are outside the scope of this framework document.

The Dynamic Delegation Discovery System (DDDS) is used to implement a binding of strings to data in order to support dynamically configured delegation systems [rfc3401]. The DDDS functions by mapping some unique string to data stored within a DDDS Database by iteratively applying string transformation rules until a terminal condition is reached. The potential then exists that a client could specify a set of known tags (e.g., RetrieveMail:Pop3) that would identify/discover the appropriate server with which it can communicate.

4.7. Differentiated Services (Diffserv)

There are a number of examples where Diffserv [rfc2474] has been deployed strictly within a domain, with no extension of service to neighboring domains. Various reasons exist for Diffserv not being widely deployed in an inter-domain context, including ones rooted in the complexity and problems in supporting the security requirements for Diffserv codepoints. An extensive discussion on Diffserv deployment is outside the scope of this document.

[Baker] presents common examples of various codepoints used for well-known applications. The document does not recommend these associations as being standard or fixed. Rather, the examples in [Baker] provide a reference point for known deployments that can act as a guide for other network administrators.

An argument can be made that Diffserv, with its existing codepoint specifications of Assured Forwarding (AF) and Expedited Forwarding (EF), goes beyond what would be needed to support ETS within a domain. By this we mean that the complexity in terms of maintenance and support of AF or EF may exceed or cause undue burden on the management resources of a domain. Given this possibility, users or network administrators may wish to determine if various queuing mechanisms, like class-based weighted fair queuing, is sufficient to support ETS flows through a domain. Note, as we stated earlier in Section 2, over-provisioning is another option to consider. We assume that if the reader is considering something like Diffserv, then it has been determined that over-provisioning is not a viable option given their individual needs or capabilities.

5. Security Considerations

Services used to offer better or best available service for a particular set of users (in the case of this document, ETS users) are prime targets for security attacks or simple misuse. Hence, administrators that choose to incorporate additional protocols/services to support ETS are strongly encouraged to consider new policies to address the added potential of security attacks. These policies, and any additional security measures, should be considered independent of any mechanism or equipment that restricts access to the administrative domain.

Determining how authorization is accomplished is an open issue. Many times the choice is a function of the service that is deployed. One example is source addresses in an access list permitting senders to the multicast group (as described in Section 4.5). Within a single domain environment, cases can be found where network administrators tend to find this approach acceptable. However, other services may

require more stringent measures that employ detailed credentials, and possibly multiple levels of access and authentication. Ease of use, deployment, scalability, and susceptibility to security breach all play a role in determining authorization schemas. The potential is that accomplishing this for only a single domain may be easier than at the inter-domain scope, if only in terms of scalability and trust.

6. Summary Comments

This document has presented a number of protocols and complementary technologies that can be used to support ETS users. Their selection is dictated by the fact that all or significant portions of the protocols can be operated and controlled within a single administrative domain. It is this reason why other protocols, like those under current development in the Next Steps in Signaling (NSIS) working group, have not been discussed.

By listing a variety of efforts in this document, we avoid taking on the role of "king maker" and at the same time indirectly point out that a variety of solutions exist in support of ETS. These solutions may involve QoS, traffic engineering, or simply protection against detrimental conditions (e.g., spikes in traffic load). Again, the choice is up to the reader.

7. Acknowledgements

Thanks to Ran Atkinson, Scott Bradner, Jon Peterson, and Kimberly King for comments and suggestions on this document.

8. References

8.1. Normative Reference

[rfc4375] Carlberg, K., "Emergency Telecommunications Services (ETS) Requirements for a Single Administrative Domain", RFC 4375, January 2006.

8.2. Informative References

[Baker] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, August 2006.

[docsis] "Data-Over-Cable Service Interface Specifications: Cable Modem to Customer Premise Equipment Interface Specification SP-CMCI-I07-020301", DOCSIS, March 2002, http://www.cablemodem.com.

[iso15802] "Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Common specifications - Part 3: Media Access Control (MAC) Bridges: Revision. This is a revision of ISO/IEC 10038: 1993, 802.1j-1992 and 802.6k-1992. It incorporates P802.11c, P802.1p and P802.12e." ISO/IEC 15802-3:1998"

[rfc1546] Partridge, C., Mendez, T., and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.

[rfc2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[rfc2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[rfc2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

[rfc2608] Guttman, E., Perkins, C., Veizades, J., and M. Day, "Service Location Protocol, Version 2", RFC 2608, June 1999.

[rfc2748] Durham, D., Ed., Boyle, J., Cohen, R., Herzog, S., Rajan, R., and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000.

[rfc2749] Herzog, S., Ed., Boyle, J., Cohen, R., Durham, D., Rajan, R., and A. Sastry, "COPS usage for RSVP", RFC 2749, January 2000.

[rfc2750] Herzog, S., "RSVP Extensions for Policy Control", RFC 2750, January 2000.

[rfc3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[rfc3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.

[rfc3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[rfc3344] Perkins, C., Ed., "IP Mobility Support for IPv4", RFC 3344, August 2002.

[rfc3084] Chan, K., Seligson, J., Durham, D., Gai, S., McCloghrie, K., Herzog, S., Reichmeyer, F., Yavatkar, R., and A. Smith, "COPS Usage for Policy Provisioning (COPS-PR)", RFC 3084, March 2001.

[rfc3401] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part One: The Comprehensive DDDS", RFC 3401 October 2002.

[rfc3535] Schoenwaelder, J., "Overview of the 2002 IAB Network Management Workshop", RFC 3535, May 2003.

[rfc3569] Bhattacharyya, S., Ed., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, July 2003.

[rfc3618] Fenner, B., Ed., and D. Meyer, Ed., "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.

[rfc4190] Carlberg, K., Brown, I., and C. Beard, "Framework for Supporting Emergency Telecommunications Service (ETS) in IP Telephony", RFC 4190, November 2005.

[rfc4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[rfc4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

Author’s Address

Ken Carlberg
G11
123a Versailles Circle
Baltimore, MD
USA

EMail: carlberg@g11.org.uk

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.