Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 21, 2008

O. Bonaventure
D. Saucez
B. Donnet
Universite catholique de Louvain
February 18, 2008

The case for an informed path selection service
draft-bonaventure-informed-path-selection-00.txt

Status of this Memo

Copyright Notice

Abstract

With today's peer-to-peer applications, more and more content is available from multiple sources. In tomorrow's Internet hosts will have multiple paths to reach one destination host with the deployment of dual-stack IPv4/IPv6 hosts, but also with new techniques such as shim6 or other locator/identifier mechanisms being discussed within the IRTF RRG. All these hosts will need to rank paths in order to select the best paths to reach a given destination/content. In this

draft, we propose an informed path selection service that would be queried by hosts and would rank paths based on policies and performance metrics defined by the network operator to meet his traffic engineering objectives. A companion document describes a protocol that implements this service.

Table of Contents

1. Introduction

The current Internet is based on several assumptions that have driven the development of most Internet protocols and mechanisms. A first assumption is that (usually) one address is associated to each host. Also, the forwarding of packets is exclusively based on the destination address. For this reason, there is usually a single path between one source (or client) and one destination (or server). Finally, the Internet was designed with the client-server model in mind assuming that many clients receive information from (a smaller number of) servers.

During the last years, these assumptions have been severely challenged :

- The client-server model does not correspond to the current operation of many applications. First, large servers are usually replicated and different types of content distribution networks are used to efficiently distribute content. Second, the proliferation of peer-to-peer applications implies that most clients also act as server. This is creating several problems in many ISP networks [1]. The client-server asymmetry does not hold anymore as earlier.

- Due to the transition from IPv4 to IPv6 many hosts will be dual-stack for the foreseeable future [2]. Furthermore, measurements show that IPv4 and IPv6 do not provide the same performance [3], even for a single source-destination pair. This implies that to reach a destination supporting both IPv4 and IPv6, a source will need to select the utilization of IPv4 or IPv6.

- Host based Multihoming techniques such as [4] are emerging. These techniques assume that each host of a multihomed site will have several IPv6 addresses (e.g., one per provider).

- Several locator/identifier separation protocols [5] [6] being discussed within the IRTF Routing Research Group allow one identifier to be reachable via multiple locators.

A consequence of the deployment of these new techniques is that the number of end-to-end paths that are available to reach a given destination/content will grow. Several studies and practical experience show that resilience of the Internet increases with the number of paths [7][8] since if one path fails, it is likely that the other paths will continue to work provided that they are sufficiently disjoint. Also, the availability of multiple paths may allow a better use of the Internet infrastructure by providing better paths in terms of delay, bandwidth, and congestion compared to the unique current IPv4 paths. This has been shown by several measurements studies [8][9].

However, to obtain these benefits, the hosts (or the routers in some of the proposals being discussed within the IRTF RRG), will need to be able to accurately select the best path to use to reach a given (set of) destination(s). Several solutions have been proposed to allow P2P applications to rank some paths over others [10] [11] [12]. However, relying on proprietary solutions implies a duplication of efforts (e.g. different peer-to-peer applications may use different techniques and perform their own measurements). Also, the existing solutions such as the static source address selection mechanism defined in [13] are static.

In this document, we propose an informed path selection service that is able to rank paths based on policy and performance criteria. A protocol to implement this service is described in a companion document [14].

This document is organized as follows. First, we provide a high-level description of the proposed service in Section 2. Then, to illustrate the benefits of such a service, we recall in Section 3 three issues for multihomed networks expressed by J. Schiller in [15]. In Section 4 we explain the limitations of existing (i.e., BGP) and proposed techniques (i.e., shim6 and LISP) when solving these case studies. In Section 5 we discuss several possible applications of the informed path selection service. Finally, we compare the informed path selection service with related work in Section 6.

2. The informed path selection service

The informed path selection service is a distributed request-response service that allows to rank paths. This service is typically supported inside a domain. It can benefit from cooperation between domains but does not require it.

```
                                BGP, OSPF/ISIS    Measurements
                                  ||                  ||
                                  ||                  ||
                                  \/                  \/
                                +-------------------------+
                                |                         |
                                |    Informed Path        |
                                |      Selection          |
                                |       Service           |
              +--------+   request  |                         |
              | client | -->----    |                         |
              |        |   -----<--- |                         |
              +--------+  response  +-------------------------+
                                         /\
                                         ||
                                      Policies
```

Informed path selection service

Figure 1

The informed path selection service is used to decide the best path(s) among a set of candidate paths. It can be queried by a host having multiple addresses, a LISP router or other entities that need to rank paths such as peer-to-peer applications, content distribution networks, dual-stack hosts, ... The informed path selection service is based on a request/response mechanism and the path ranking may depend on several factors including :

- Routing information (e.g., BGP, OSPF/ISIS) that allow the informed path selection service to compare different paths based on routing metrics (e.g. BGP local preference, BGP AS-Path length, IGP path length, ...).

- Active or passive measurements (e.g., delay, bandwidth, loss, ...) that allow the informed path selection to compare different paths based on quantitative performance metrics.

- Policies configured by the network administrator that indicate preferences for some paths over others.

A request will contain the following information :

- one or more source addresses (or prefixes),

- one or more destination addresses (or prefixes).

Upon reception of a request, the informed path selection service builds a list of all the possible paths between the source(s) and the destination(s). Then, it removes from consideration the paths that are invalid due to routing (e.g., one destination is not reachable from a given source address) or policies. These remaining paths are ranked and the reply contains the following information :

- the best path (source address, destination prefix),
- the second best path (source address, destination prefix),
- ...
- the Nth best path (source address, destination prefix),
- the lifetime for the ranked paths.

As indicated above, the number of paths returned by the path selection service may be lower than the total number of possible paths, e.g., because some paths are not usable due to policy reasons or because some destinations are not reachable by using some source addresses.

For scalability reasons and based on the experience in developing the NAROS protocol [16], the informed path selection service uses two mechanisms to allow the client to use the same path for several flows. First, an ordered list of paths is valid for some time and the client is encouraged to cache the ordered list for the lifetime indicated in the response. Second, the response may contain paths that are composed of a source and a destination prefix instead of addresses. This choice is motivated by the fact that all the IP addresses that belong to the same prefix are usually covered by the same policies and have similar performance. Simulations performed earlier showed that these two mechanisms allow to significantly reduce the number of requests sent to an informed path selection service by a given host [8].

3. Issues with multihoming

To illustrate the benefits of an informed path selection service and compare it with existing techniques, we first summarize the concerns raised by J. Schiller on multihoming in [15]. We focus on the three main case study described in [15].

3.1. Case 1 : Primary/Backup

The first issue mentioned in [15] is the classical case of a multihomed network using one primary provider and another as backup.

This case is illustrated in Figure 2 where the multihomed customer is attached to UUNet/MCI and ATT. In this example, traffic engineering objectives of the multihomed customer are :

- o All outgoing packets must be sent via UUNet/MCI when this link is active. Otherwise, the packets must be sent via ATT.

- o All incoming packets must be received via UUNet/MCI when this link is active. Otherwise, the packets must be received via ATT.

```
                 +---------------+       +---------------+
                 |   UUNET/MCI   +-~-~-~-+      AT&T     |
                 +------:--------+       +-------:-------+
                        \                       //
                         \                     //
                          \                   //
                           \                 //
                            \               //
                             \             //
                              \           //
                               \         //
                                \       //
                                 \     //
                             +----:----:-----+
                             |   Multihomed  |
                             |    customer   |
                             | 63.63.62.0/23 |
                             +---------------+

                 --- primary link to UUNET
                 === backup link to AT&T
                 -~- Internet
```

Example of a Primary/Backup implementation

Figure 2

This problem has been generalized in [15] as follows : "It is required to have the ability to set a link or a set of links as primary and some other as backup links. Such that all the traffic is carried by the primary links while backup links are used only while primary links become unavailable."

3.2. Case 2 : Load Sharing

The second case mentioned in [15] is load sharing across links from different providers. This problem is illustrated in Figure 3. In this example, the high-level objectives of the multihomed customer can be specified as :

- o The same amount of outgoing packets should be sent via the UUNet/ MCI and ATT links.
- o The same amount of incoming packets should be received via UUNet/ MCI and ATT links.

```
                              Load sharing

              +---------------+       +---------------+
              |   UUNET/MCI   +-~-~-~-+      AT&T     |
              +------:--------+       +-------:-------+
                      \                       //
                       \                     //
                        \                   //
                         \                 //
                          \               //
                           \             //
                            \           //
                             \         //
                              \       //
                               \     //
                         +----:----:-----+
                         |   Multihomed  |
                         |    customer   |
                         | 63.63.62.0/23 |
                         +---------------+

              --- link to UUNET
              === link to AT&T
              -~- Internet
```

Figure 3

This problem can, of course, be generalized by considering more than two links/providers and also by requiring unequal load sharing among the different links (e.g. based on link cost, link capacity, ...).

3.3. Case 3 : Best Path

The third case mentioned in [15] is that it should be possible for the multihomed customer to have ways to decide whether the paths available via one provider are better than the paths available via the other provider and use the best paths. The high-level objectives of the multihomed customer in this case become

- For each destination, send the outgoing packets via the provider that has the best path to reach this destination.

- For each source, receive the incoming packets via the provider that is on the best path from this source.

This problem can be generalized to cover more than two links and providers.

4. Existing multihoming solutions

In this section, we evaluate how three technical solutions that are used today or are being discussed within the IETF/IRTF are able to meet the objectives mentioned above. We first start with BGP-based multihoming as described in [15], then discuss shim6 [4] and finally LISP [5].

4.1. BGP-based multihoming

BGP-based multihoming is a common and widely deployed technique that allows a multihomed network to be attached to different providers. It is used by existing IPv4 and IPv6 deployments.

4.1.1. Case 1 : Primary/Backup

With BGP-based traffic engineering, the common techniques to implement primary/backup links are the following [15]:

- Set a higher MED for backup links from the same AS.

- Set a lower local-preference for backups links of different ASes.

- Set a higher weight for static default routes on backup links.

The main issue with these BGP-based solutions is that a prefix must be allocated to each multihomed customer. Furthermore, this prefix is advertised in the BGP routing tables and thus contributes to the growth of these routing tables [17].

4.1.2. Case 2 : Load Sharing

To solve the load sharing case, several techniques can be used on BGP routers. The following are presented in [Schiller-TE] [15]

- Divide IP space and more specific routes announces over the different links.

o Modify the MED or local preferences of inbound links.

o Modify IGP metrics to move hosts closer to a given exit point.

o Manipulate equal cost static default routes.

Figure 4 from [15] shows how BGP can be used to solve the load sharing problem by dividing the IP space of the multihomed customer and sending more specific routes. This solution has several drawbacks. First, it contributes to the growth of the BGP routing tables by requiring each multihomed customer to advertise more than one prefix [17]. Second, the solution is far from perfect and assumes that the two /23 more specific prefixes carry almost the same amount of packets. If this is not the case or if the amount of packets changes with time, then the more specific prefixes that need to be advertised also need to change with time. This is not desirable as a multihomed customer willing to move some packets from one link to another would need to send BGP updates that would change over time.

```
              +---------------+          +---------------+
              |   UUNET/MCI   +-~-~-~-~-~+      AT&T     |
              +------:--------+          +-------:-------+
                      \                          //
                       \                        //
                        \                      //
                         \                    //
             advertise 63.63.62.0/23   advertise 63.63.62.0/23
             advertise 63.63.62.0/24   advertise 63.63.63.0/24
                            \             //
                             \           //
                              \         //
                               \       //
                 receive default\     //receive default
                            +----:----:-----+
                            |   Multihomed  |
                            |    customer   |
                            | 63.63.62.0/23 |
                            +---------------+

             --- primary link to UUNET
             === primary link to AT&T
             -~- Internet
```

Example of a load sharing implementation

Figure 4

4.1.3. Case 3 : Best Path

With BGP-based multihoming, several techniques can be used to select the best path based on different definitions of best. They all require the multihomed customer to receive the full BGP routing tables from its providers and run BGP. A drawback of this solution is that the definition of "best path" either depends on the limited BGP attributes or must be tuned manually. Measurements have shown that there is not a strong correlation between the length of the AS Path carried in BGP messages and the performance of path measured in terms of delay or bandwidth [18].

- o Best path is selected according to the BGP Decision process depicted in [19] section 9.1.

- o Traffic is controlled by the BGP path selection algorithm of the source of the traffic.

- o Manual changes can be used to move traffic from over-loaded links to under-loaded links.

4.2. Shim6 host-based multihoming

The shim6 host-based multihoming technique is based on the assumption that multihomed hosts will use one IPv6 address per provider. It is expected that most multihomed customers will use PA addresses. In this case, the multihomed customer does not need to advertise any prefix with BGP. The basic network scenario with shim6 is depicted in Figure 5.

```
+---------------+         +---------------+
|   UUNET/MCI   +-~-~-~-~-+      AT&T     |
|               |         |               |
|   2001::/48   |         |    2002::/48  |
+------:--------+         +-------:-------+
        \                         //
         \                       //
          \                    //
           \                  //
            \                //
             \              //
              \            //
               \          //
            +---:--------:---+
            |   Multihomed   |
            |    customer    |
            |                |
            |   2001::1:A    |
            |   2002::A:1    |
            +----------------+

--- primary link to UUNET
=== backup link to AT&T
-~- Internet

Example of a Primary/Backup implementation
```

Figure 5

4.2.1. Case 1 : Primary/Backup

With the current IPv6 and shim6 specifications, a primary/backup implementation can be supported by using the default address selection specified in [13]. This specification defines a table that is used by each host to select the source address that it will use to reach a given destination based on the type of address, the prefix length and preferences that can be defined by the system administrators. This solution allows all hosts to prefer once source address over the other. However, no protocol has been defined to allow a system administrator to distribute the current preferences to its hosts. This implies that the preference is rather static.

4.2.2. Case 2 : Load Sharing

Concerning load sharing, the current shim6 specification [4] can be configured by using the preferences of the source address selection mechanism to prefer one link over the other. As with BGP-based multihoming, this solution is static, it is difficult to dynamically change the source address selection preferences of the hosts to follow the evolution of the traffic patterns.

4.2.3. Case 3 : Best Path

The current shim6 specification does not expect that the hosts will select the source and destination addresses for a shim6 session based on performance metrics but does not preclude it. An unrealistic option would be to add a BGP and IGP routing table on each host to allow them to select the best (source address,destination address) pair based on BGP metrics. Additional information about operator's concerns with shim6 may be found in [20].

4.3. Dual stack IPv4/IPv6

Several mechanisms have been proposed to ease the transition from the current IPv4 Internet towards an IPv6 Internet. As of today, nobody expects that the IPv6 Internet will completely replace the IPv4 Internet quickly and that both Internets will coexist for several years or more. For the foreseeable future, many networks will be attached to both IPv6 and IPv4 providers. When considering the three case studies, the dual stack IPv4/IPv6 hosts have several problems as shim6 hosts discussed in the previous section. The only difference is that a host cannot switch from using IPv4 to using IPv6 for an established flow (i.e., if IPv4 connectivity becomes broken but IPv6 connectivity remains active).

4.4. LISP and multihoming issues

The Locator/Identifier Separation Protocol (LISP) [5] is currently being discussed within the IRTF Routing Research Group as one of the possible alternatives to achieve a better scaling of the Internet architecture. LISP distinguishes between identifiers and locators. The identifiers are used to identify endhosts. The locators are assigned to ingress routers that implement the LISP tunneling scheme. When an endhost needs to contact a remote endhosts, it sends a packet with its own identifier as source address and the identifier of the remote host as destination address. This packet will be intercepted by the first LISP router on its path. This router uses a mapping system to query the locator that allows to reach the (identifier of the) remote host and encapsulates the packet before sending it to the locator associated to the remote host.

```
             +----------------+          +----------------+
             |   UUNET/MCI    +-~-~-~-~-~+      AT&T      |
             | 210.0.0.0/8    |          |   11.0.0.0/8   |
             +-------:--------+          +--------:-------+
                      \                          //
                       \                        //
                        \                      //
                         \                    //
                          \                  //
                           \                //
                            \              //
                             \            //
                           +----:----:-----+
                           |   Multihomed  |
                           |    customer   |
                           |    Locators   |
                           | 210.1.2.0/29  |
                           | 11.200.2.0/28 |
                           +---------------+

            --- primary link to UUNET
            === backup link to AT&T
            -~- Internet
```

Example of LISP scenario

Figure 6

4.4.1. Case 1 : Primary/Backup

With the current LISP specification, the primary/backup case can be covered by considering the priority that is associated to an EID/locator mapping. A LISP router will prefer the locator having the highest priority. This allows each LISP router to select the best locator to reach a given destination.

4.4.2. Case 2 : Load Sharing

In addition to the priority mentioned above, LISP also associates a weight to each locator. When several locator have the same priority, then load sharing should be performed among the different locators based on their weight.

4.4.3. Case 3 : Best Path

If the LISP routers of the multihomed site run BGP, they can use the BGP decision process to rank some routes over others. However, as explained earlier, the correlation between BGP attributes such as the length of the AS Path and the performance of interdomain paths is weak.

5. Application of the informed path selection service

In this section, we briefly discuss how the informed path selection service could improve the performance of multihomed networks by considering our three case studies. As an example, we consider that the informed path selection service is queried by hosts, but the same result would apply for LISP routers.

5.1. Case 1 : Primary/Backup

By using the informed path selection service, the primary/backup case can be easily solved. Upon reception of a request, the server simply needs to always place the prefixes that correspond to the primary link at the top of the list and possibly remove the prefixes associated to the backup link from the reply. When the primary link fails, the server updates its ranking to allow the hosts to use the backup link instead.

5.2. Case 2 : Load Sharing

The load sharing case can be naturally solved by using an informed path selection service. Indeed, the service could easily track the load on the different links and dynamically change its replies based on the link load. The NAROS protocol [16], proposed in the early days of IPv6 multihoming, was designed to solve this problem and the evaluation showed that it worked well [8].

5.3. Case 3 : Best Path

The informed path selection service brings new benefits for the Best path case as it allows the server to base its ordering on active measurements to assess the performance of paths by considering metrics such as delay or bandwidth. The informed path selection service is not restricted to the BGP information as in the current BGP-based multihoming techniques.

5.4. Other applications of the informed path selection service

The informed path selection service is not limited to multihomed networks. It can be used in any environment where several paths need to be ranked based on policies and/or performance.

The peer to peer applications are clear candidate users for such a service. Some peer-to-peer applications already rely on heuristics to prefer some sources over others. A standardized path selection service would allow several peer-to-peer applications to share the same measurements. Furthermore, an ISP or campus network running the informed path selection service could influence providers used by the packets sent/received by the hosts of its networks.

The informed path selection service could be associated to a DNS resolver or server. When a DNS resolver receives a DNS reply containing several addresses for the same name, it could rank them and return a ranked DNS response. A DNS server implementing [21] could contact the informed path selection service to update dynamically the SRV RR of its local servers.

The informed path selection service could also be useful for multihomed VoIP gateways that need to select the best VoIP gateway to forward a voice call.

6. Related work

Several solutions have been proposed to improve the performance of end-to-end paths. A first approach was proposed with the RSVP signaling protocol [22] and the Integrated Services Architecture [23]. RSVP allows to reserve resources on all routers along and end-to-end path but does not allow a host to prefer one path over another. Other signaling protocols have been or are being proposed to install and maintain state on some intermediate nodes [24]. Our proposed path selection service follows the end-to-end principle [25] and does not create any state in intermediate nodes.

Due to scalability concerns, the Integrated Services Architecture has not been widely deployed. Differentiated Services [26] were introduced as a more scalable solution based on packet marking. Differentiated Services does not by itself allows hosts to prefer some paths over others. However, recent extensions to link state routing protocols or the utilization of MPLS allow network operators to provision different paths for different classes of services. Our proposed path selection service allows the client to also indicate a DSCP in the request to support hosts and applications that are using non best-effort service. Although it does not require Differentiated services, it can easily cooperate with it.

Several researchers have proposed solutions to similar problems. For example, [27] proposed a mechanism where the source prefix of shim6 data packets is rewritten by the site routers. The proposed informed path selection service does not require routers to change source prefixes. [12] proposed an oracle service that would be configured by the network operator and queried by peer-to-peer applications. The oracle could be one of the ways to implement a path selection service. Other mechanisms have been proposed specifically for peer-to-peer applications [28] [10] [11].

7. Security Considerations

By ranking paths, the informed path selection service influences the path that hosts will use to send packets to some destinations. By controlling the informed path selection service, an attacker diverts packets through a path that he controls to create man-in-the middle attacks or divert packets over an overload path to increase congestion. These problems are similar to the security issues with DNS resolver since an attacker who controls a DNS resolver could obtain similar results. To mitigate these risks, it should be possible for the clients that are using the informed path selection service to authenticate the responses received from a server.

8. Conclusion

In this document, we have proposed an informed path selection service that is able to rank paths based on policies or performance criteria. A companion document [14] proposes a protocol to support this service.

9. Acknowledgements

This work was partially supported by the European Commission within the IST AGAVE and IST ONELAB projects.

10. Normative References

[1] Karagiannis, T., Rodriguez, P., and K. Papagiannaki, "Should Internet Service Providers Fear Peer-Assisted Content Distribution?", Internet Measurement Conference 2005, October 2005.

[2] Cho, K., Luckie, M., and B. Huffaker, "Identifying IPv6 network problems in the dual-stack world", ACM SIGCOMM Workshop on Network Troubleshooting 283 - 288, September 2004.

[3] Zhou, X., Jacobsson, M., Uijterwaal, H., and P. Van Mieghem, "IPv6 delay and loss performance evolution", International Journal of Communication Systems DOI: 10.1002/dac.916, 2007.

[4] Bagnulo, M. and E. Nordmark, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", draft-ietf-shim6-proto-09 (work in progress), November 2007.

[5] Farinacci, D., "Locator/ID Separation Protocol (LISP)", draft-farinacci-lisp-05 (work in progress), November 2007.

[6] Vogt, C., "Six/One: A Solution for Routing and Addressing in IPv6", draft-vogt-rrg-six-one-01 (work in progress), November 2007.

[7] Andersen, D., Balakrishnan, H., Kaashoek, F., and R. Morris, "Resilient Overlay Networks", Proc. 18th ACM SOSP Banff, Canada, October 2001.

[8] de Launois, C., "Unleashing traffic engineering for IPv6 multihomed sites", PhD thesis available from http://inl.info.ucl.ac.be/delaunoi, October 2005.

[9] Akella, A., Maggs, B., Seshan, S., Shaikh, A., and R. Sitaraman, "A measurement-based analysis of multihoming", Proc. SIGCOMM (Karlsruhe, Germany, August 2003.

[10] Xie, H., Krishnamurthy, A., Yang, R., and A. Silberschatz, "P4P: Proactive Provider Participation for P2P", Yale Computer Science YALE/DCS/TR1377 , March 2007.

[11] Dabek, F., Cox, R., Kaashoek, F., and R. Morris, "Vivaldi: a decentralized network coordinate system", Proc. SIGCOMM 2004 15-26, August 2004.

[12] Aggarwal, V., Feldmann, A., and C. Scheideler, "Can ISPs and P2P systems co-operate for improved performance?", ACM SIGCOMM Computer Communications Review Volume 37, Number 3, July 2007.

[13] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[14] Saucez, D., Donnet, B., and O. Bonaventure, "IDIPS : ISP-Driven Informed Path Selection", Internet-Draft

draft-saucez-idips-00, February 2008.

[15] Schiller, J., "Inter-AS Traffic Engineering Case Studies as Requirements for IPv6 Multihoming Solutions", NANOG 35, May 2005.

[16] Launois, C. and O. Bonaventure, "NAROS : Host-Centric IPv6 Multihoming with Traffic Engineering", draft-de-launois-multi6-naros-00 (work in progress), May 2003.

[17] Meyer, D., "Report from the IAB Workshop on Routing and Addressing", draft-iab-raws-report-02 (work in progress), April 2007.

[18] Huffaker, B., Fomenkov, M., Plummer, D., and k. Claffy, "Distance Metrics in the Internet", IEEE International Telecommunications Symposium , September 2002.

[19] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

[20] Schiller, J., "Shim6 : network operator concerns", IAB Workshop , October 2005.

[21] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.

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

[23] Braden, B., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.

[24] Manner, J. and X. Fu, "Analysis of Existing Quality-of-Service Signaling Protocols", RFC 4094, May 2005.

[25] Saltzer, J., Reed, D., and David. Clark, "End-to-end arguments in system design", ACM Transactions on Computer Systems Vol 2, N 4 (November 1984) pages 277-288, November 1984.

[26] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[27] Bagnulo, M., Garcia-Martinez, A., and A. Azcorra, "BGP-like TE Capabilities for SHIM6", EUROMICRO 2006, September 2006.

[28] Ledlie, J., Gardner,, P., and M. Seltzer, "Network Coordinates in the Wild", Proceedings of NSDI 2007 , April 2007.

Authors' Addresses

Olivier Bonaventure
Universite catholique de Louvain
Place Sainte Barbe 2
Louvain-la-Neuve,   1348
Belgium

URI: http://inl.info.ucl.ac.be

Damien Saucez
Universite catholique de Louvain
Place Sainte Barbe 2
Louvain-la-Neuve,   1348
Belgium

Email: damien.saucez@uclouvain.be
URI: http://inl.info.ucl.ac.be

Benoit Donnet
Universite catholique de Louvain
Place Sainte Barbe 2
Louvain-la-Neuve,   1348
Belgium

Email: benoit.donnet@uclouvain.be
URI: http://inl.info.ucl.ac.be

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).