

BGP: The Next Best Thing Since Sliced Bread?

Petri Miettinen

Helsinki University of Technology

Telecommunications Software And Multimedia Laboratory

Petri.Miettinen@hut.fi

1 Abstract

Border Gateway Protocol, BGP, is de-facto standard for routing today in Internet. Current version of BGP is 4 and it supports internal and external routing in an Autonomous System.

The evolution of BGP has converted it from a routing protocol into a multiprotocol transport mechanism. This evolution has made IETF worried about the stability of global Internet routing. Adding new applications into existing protocol will have consequences, at least in the form of additional complexity of the system. In this paper we will take a look at BGP and its recent development. We will also study implementations of VPN, VPWS and VPLS in BGP. It seems that VPN fits quite well into BGP architecture but the situations of VPLS and VPWS are still developing. It might be that VPLS and VPWS have some requirements too specific for BGP data distribution model.

This paper is a summary from several papers written by the experts of the area and somewhat a personal opinion about the extension of BGP. It seems that adding new application into BGP has to have a solid reason. We should not forget the main job of BGP, route information in Internet. Everything else is secondary.

KEYWORDS: BGP, Border Gateway Protocol, Multiprotocol Extensions of BGP, BGP Extended Communities Attribute, Flexible BGP Communities, RIFT, VPLS, VPWS, SPT, GPT

2 Introduction

2.1 BGP

BGP, or Border Gateway Protocol, is an intelligent routing protocol used widely today in Internet. It is de facto standard of exterior routing between *Autonomous Systems* (ASs). An Autonomous System is a unit of router policy which is controlled by a single network administrator. A unit can be a single network or a group of networks. It has been extended to perform other tasks than routing too.

2.2 Problem Statement

The problem this paper tries to solve is to analyze whether BGP should be used to do other tasks besides routing. More specifically, we will study Virtual Private Networks (VPN), Virtual Private LAN Services (VPLS) and Virtual Private

Wire Services (VPWS). These three services are implemented to BGP and we will study possible benefits and/or drawbacks of these implementations.

To summarize the problem, we could ask whether “BGP is the next best thing after sliced bread?”. Is it really the general solution to common problems in current network environment?

2.3 Scope of this Work

The intention of this paper is to study BGP routing protocol and its added functionalities. The study is made purely on theoretical basis using material written by the experts of the area.

3 Introduction to BGP

3.1 What is BGP?

Border Gateway Protocol is current de-facto standard of routing. BGP was designed to substitute Exterior Gateway Protocol (EGP) which had some serious limitations. EGP was created around backbone centered tree which is not the actual case anymore in Internet. Neither did EGP accommodate to the growth of the Internet. The current version of BGP is 4 (also written as BGP-4) and is hereafter referred to as BGP if not stated otherwise. BGP-4 was introduced in [1] in March 1995 and it is required for Classless Inter-Domain Routing (CIDR).

3.2 Technical Details

BGP is a protocol running over TCP (Transmission Control Protocol) and uses port number 179. The use of TCP eliminates the need of retransmission and acknowledgement in BGP protocol. BGP supports two types of routing: internal and external routing. External routing (EBGP) means exchanges between different ASs (Autonomous Systems) and internal routing (IBGP) is used in peers within an AS. Autonomous System means a collection of networks with the same routing policy. They use single routing protocol and are probably under one ownership and administrative control. BGP speakers are called as *peers*.

BGP is classified as *path vector protocol* (see [2]). “A path vector protocol defines a route to a destination as a pairing between the destination and the attributes of the path to the destination”[2]. BGP peer learns paths from internal and

external speakers. It picks up the best path and updates the forwarding table.

3.3 Message Types

The information between BGP peers is exchanged by sending messages. The maximum size of a message is 4096 octets. BGP can send 4 types of messages: OPEN, UPDATE, KEEPALIVE and NOTIFICATION.[1]

OPEN message is sent right after the connection is established. Both ends of the connection send it and reply to it by sending KEEPALIVE. After the confirmation made by KEEPALIVE, other messages can be sent.[1]

UPDATE message is used to send actual routing information. UPDATE can be either to advertise one route or withdraw multiple routes. UPDATE message can include several attributes which describe the route. The attributes are discussed in the following chapter.[1]

Since BGP does not use any transport protocol keep-alive mechanism, the peers send periodically KEEPALIVE messages to determine whether other end is alive or not. Both ends have Hold Timer value which is the maximum time of inactivity before the other end is assumed to be dead.[1]

NOTIFICATION message is sent when an error condition occurs. After this message the connection is immediately closed. The error may occur for example in OPEN or UPDATE message, or Hold Timer can expire.[1]

3.4 Attributes

BGP contains attributes which describe paths and helps to choose the best route among them. The attributes described are used in path selection process as described in [3] These attributes are: Weight, Local Preference, Multi-Exit Discriminator, Origin, AS_path, Next_hop and Community

WEIGHT is a Cisco-defined attribute. It is not advertised to neighbouring routers. If there are more than one routes to the destination, the one with the highest weight will be chosen.

LOCAL PREFERENCE is used inside an AS to prefer certain exit point from the AS. If there are more than one exit point inside the AS, the one with highest local preference will be chosen.

Multi-Exit Discriminator (MED) is a suggestion to an external AS regarding the entry point to the originating AS. It is a suggestion since the external AS may be using other attributes when selecting the route to the AS.

ORIGIN attribute tells where the specific route was learned from. There are 3 different values for this attribute: IGP means the route is interior to the AS, EGP means the route was learned via EGP and Incomplete implies that the origin of the route is unknown.

AS_PATH attribute was created to prevent routing loops. When a route advertisement passes an AS, the AS adds its own AS number to ordered list of AS numbers where the advertisement has passed.

NEXT_HOP attribute is in EBGp the IP address that is used to reach the advertising router. Inside an AS, the next-hop attribute is carried without modifications.

COMMUNITY attribute can be used to group communities, or destinations. These groups can be used to apply routing decisions differently. Predefined community attributes are NO_EXPORT, NO_ADVERTISE and NO_EXPORT_SUBCONFED. NO_EXPORT means the route is not advertised to BGP peers outside a confederation (confederations explained in [4]), NO_ADVERTISE means the route is not advertised to any peer and NO_EXPORT_SUBCONFED that the route must not be advertised to any EBGp peer.[5]

4 Evolution of BGP

From its beginning in 1988, BGP has evolved around real problems of the Internet. From BGP-1 the protocol has been converted from "simple" routing protocol into multiprotocol transport mechanism. The evolution has happened so fast that in 2000-2001 Internet Engineering Task Force (IETF) became concerned about the BGP infrastructure[6]. This chapter will study some of the newest additions to BGP protocol.

4.1 Multiprotocol Extension for BGP

BGP is fundamentally linked to IPv4. There are three pieces of information which contain IPv4 address: NEXT_HOP attribute, AGGREGATOR attribute and NLRI (Network Layer Reachability Information) field. To support multiple protocols, two new attributes were created. This also enables backward compatibility since the attributes were defined as optional and non-transitive. It was assumed that the BGP speaker has a valid IPv4 address which is used for example in AGGREGATOR attribute. The attributes created were *Multiprotocol Reachable Network Layer Reachability Information (MP_REACH_NLRI)* and *Multiprotocol Unreachable NLRI (MP_UNREACH_NLRI)*. These attributes contain address family identifier and makes possible the use of different protocols than IPv4. The attributes are defined in [7] which was published June 2000.

The existence of the new attributes also enabled BGP to carry a variety of data types and their signalings. The combination of data types and signaling is often called an 'application'. Examples of new applications adopted to BGP are support for IPv6 (see [8]), Flow Specification Rules, IP VPNs, auto-discovery mechanism for Virtual Private Networks, Virtual Private LAN Services (VPLS) and Virtual Private Wire Services (VPWS). Some of these, like IPv6, are clearly extensions used for routing purposes but applications like VPN, VPLS and VPWS are not so clear cases. These are also the ones studied in this paper.

4.2 Internet Draft of BGP-4

The newest draft of BGP-4[9] was published in November 2003. It has mostly clarifications and deprecations when compared to [1]. One of the obvious improvement suggestions is to allow a BGP speaker to add more than one instance of its own AS to the AS_PATH attribute. This can be used in inter-AS traffic engineering. The other improvement is that

the implementations of BGP must support TCP MD5 [10] for authentication.

4.3 Multisession BGP

Most BGP implementations allow only one ESTABLISHED connection to be open between two BGP peers (two IP endpoints). Multiprotocol Extensions for BGP allow multiple NLRI families to be transported in BGP. However, one malformed message can terminate the connection between two peers. This malformed message may only be defective in its application context and not the whole connection. The purpose of multisession BGP is to allow to create several transport sessions between two peers so that they are distinct to each other.[11]

Multisession BGP offers backward compatibility so that it can communicate with the peers without the functionality. Basically the functionality is then the same as in standard BGP defined in [1]. BGP speaker with multisession capability should always advertise it in its OPEN message. If the receiving peer also supports multisession, the connection should be changed to 'multisession' mode. This should happen even if the connection is already open meaning the connection has to be closed and re-opened.

4.4 BGP Extended Communities Attribute

The extended communities attribute introduces two important enhancements to the existing Community attribute. It has an extended range which means it can be quite safely be used in applications without fear of overlapping. Extended Community Attribute has also a Type field so the communities can be structured within an application without grouping them explicitly. This has several advantages, for example making possible to divide part of the communities to be transitive and the other part non-transitive.

There are some defined BGP Extended Community Attributes. Two- and Four-octet AS specific extended communities, IPv4 address specific extended community, Opaque Extended Community, Route Target Community, Route Origin Community and Link Bandwidth Community. Link Bandwidth Community is a good example of useful extension to BGP protocol as it describes quite accurately the capability of the link to the next hop. Its use is the same as previously introduced Cisco-defined WEIGHT attribute.

BGP Extended Communities Attribute is still in the draft phase as this is written.[12]

4.5 Flexible BGP Communities

Flexible BGP Communities are built on the experience based on standard BGP community and BGP Extended Communities. Flexible BGP Communities are called the third generation of BGP Community, the first one being documented in [5] and the second, extended community in [12]. The specification is in the phase of draft.

Flexible BGP Communities have many important enhancements over standard Community and Extended Community attributes. It supports IPv6, more efficient encoding

of data, clean support of future data field structures and interpretations, support for locally defined community structures and easy extensibility for a future applications.[13]

Example of the use of Flexible BGP Community is BGP Proxy Community Community which is described in [14]. It is in the state of proposition (draft) and its use is to send another community to an external AS from certain AS.

5 Tools to analyze BGP protocol for an application applicability

There are some theoretical tools to analyze BGP and its applicability for certain applications. These tools are introduced in [15] and this section is mainly based on that text. The tools include two models, GPT and SPT, and three helping key terms: Risk, Interference and Application Fit. Two first of these terms can be analyzed just using a model and the last one is used when a certain application is analyzed. The tools are meant for application developers and network operators to take into concern before implementing an application or taking it into use.

5.1 Architectural models

There are two models which are used to analyze the whether BGP should and could be used to carry information of certain application. Two basic questions they are trying to answer are: does the application fit to BGP and what are the effects on the routing?

5.1.1 General Purpose Transport Infrastructure (GPT) Model

GPT models BGP as a general application transport mechanism. Its view point is the requirements of the attributes needed by the application; are they similar enough as the requirements of the attributes of BGP. GPT focuses on application fit and assumes risk and interference are handled. The main reason why BGP should be considered as the General Purpose Transport Protocol is that it is spread across the Internet into wide use.

5.1.2 Special Purpose Transport Infrastructure (SPT) Model

SPT sees BGP as a mechanism designed specifically to transport routing information. It, on the contrary of GPT, focuses on risk and interference. The risk is adding new features into BGP and making it possibly more unstable. The interference exposes the increased complexity after adding an application into BGP and possible destabilization of the global routing. This model can naturally be justified by the fact that any code, even unused, can make the system more unstable.

5.2 Risk, Interference, and Application Fit (RIFT)

Risk, Interference and Application Fit are tools to analyze the applicability of BGP in use of other applications. They

trade-offs between the stability of the routing system and the desire to deploy new services and reduce costs by reusing existing protocol (BGP).

5.2.1 Risk: Software Engineering

Risk models the generic software issues on a given implementation. It is the trade-off between extending an existing application and creating a new one.

The risk includes the effect of resource sharing. The resources of different applications should not compete. In general, applications in SPT model should not require competition and GPT model can avoid the competition using Multi-session BGP (discussed above).

Code impact means whether a new application can be written without affecting the old ones. Good programming practice, however, encourages the reuse of libraries, packages etc. Both models have to take this into consideration.

5.2.2 Interference: Protocol Specification/Dynamic Behaviour

Interference analyzes if new application interferes with the existing ones in some unexpected way. In other words, do we create dependency between the applications. Multisession BGP reduces the interference of the applications and the analysis in [15] shows that both models have similar profiles on interference.

5.2.3 Application Fit: Distribution Topology

Application fit analyzes how closely the requirements of the data to be distributed match the capabilities of the underlying mechanism, in this case BGP. Or in other words, if the application discussed is a routing protocol or not. The application fits are discussed in the following section.

6 Application fit and BGP

6.1 VPN

Virtual Private Network, VPN, is a mechanism to connect securely remote peers over insecure network. There are 2 major uses of VPN nowadays. The first one is to connect geographically separated networks into one, virtual network. The other use is to allow a connection to the network from a remote PC. The remote PC could be a laptop of a salesman or a home PC of an employee.

There are three kinds VPN technologies.

1. Trusted VPN is implemented usually inside single Internet Service Provider (ISP) so that the customer trusts the ISP to maintain the integrity of the circuits and do the best it can to avoid snooping.
2. Secure VPN offers an encryption to the traffic which is needed when VPN is implemented above public network.
3. Hybrid VPN is a combination of the two previously mentioned so that parts of the trusted VPN are secured with secure VPN.[16]

The implementation of VPN reside in either layer 2 or layer 3 in ISO OSI layer model. Most VPNs offered used to be in Layer 2 and they were secure and resistant to Denial-of-Service (DoS) attacks. Their problem is that they are not really scalable since they require virtual circuits for each connection. The implementations of VPN include IPsec (layer 3; described for example in [17]), IPsec inside of Layer 2 Transport Protocol ([18]), ATM circuits (layer 2), Frame Relay Circuits (layer 2; [19]) and *BGP/MPLS VPN* (layer 3; [20]).[16]

BGP/MPLS VPN defines VPN using BGP for distributing routes. *Multiprotocol Label Switching Protocol (MPLS)*[21] is used to forward packets over the backbone. A BGP speaker can install only one route to a given address prefix but each VPN can have its own address space. So more than one VPN can have the same address denoting to different system. Therefore BGP must have some way to handle multiple routes to a single address prefix. BGP Multiprotocol Extensions[7] make this possible since they allow carry routes from multiple "address families". The notion of "VPN-IPv4 address family" was introduced for BGP/MPLS VPN. It begins with 8-byte Route Distinguisher (RD) which can be used to separate same 4-byte IPv4 addresses following it.

Two PEs (*Provider Edge*) can share VPN-IPv4 routes using IBGP connection between them. If PE wants to send information to a PE in another AS, the PE needs to send the information to *Autonomous System Border Router (ASBR)* by means of IBGP and the ASBR redirects the information to the ASBR of the external AS using EBGP. This allows connecting different VPN sites to different Service Providers. It is notable that VPN-IPv4 routes should be accepted on EBGP connections only as a part of trusted arrangement. The routes should never be accepted from public Internet.

6.2 VPWS

Virtual Private Wire Service is a Layer 2 point-to-point service. As a point-to-point service, it does not have scalability issues like VPN and VPLS. This is the kind of Layer 2 VPN which has been in the markets over ATM and Frame Relay backbones[22]. The service is similar to BGP/MPLS VPN which was discussed above. The difference is that Layer 2 switching is used instead of Layer 3.

Whether or not VPWS fits into BGP is under much of discussion. The assertions (and their counter-assertions) can be found in detail from [15]. There are many open issues and it is more a matter of opinion which side one chooses. One of the issues are per-wire congestion control which would require regular point-to-point message exchanges (this is not the case in BGP routing). Actually the whole discussion whether or not VPWS fits into BGP goes around per-wire attributes. There are not (yet) many per-wire attributes defined but the situation might change in the future.

6.3 VPLS

Virtual Private LAN Service is also known as Transparent LAN Service or Virtual Private Switched Network Service. It is described as Layer 2 Virtual Private Network. The dif-

ference between traditional layer 2 VPNs is that the customers are connected by a multipoint network and not by point-to-point as in usual Layer 2 VPNs. VPLS glues several LANs so that for customers it appears they are using a single LAN, i.e. there exists PE-to-PE tunnels between different networks.[23]

There are two fundamental functions of VPLS: auto-discovery and signaling. Auto-discovery means the automatic process of discovery of other PEs (Provider Edge) which participate in a given VPLS. Other way is to pre-configure to the PE the identities of other PEs. The latter has serious configuration overhead when the topology of VPLS changes. Signaling refers to setup and teardown of the pseudowires between the connectors.[23]

For auto-discovery, VPLS uses BGP's extensions, MP_REACH_NLRI and MP_UNREACH_NLRI, to identify the members of the VPLS. It also uses Extended Route Target Attribute as also proposed in [24]. In signaling, BGP's NLRI field is used with new Address Family Identifier (AFI) and Subsequent AFI (see [7]).

The focus of VPLS is to distribute the data inside the VPLS and not to any external agent. VPLS in itself does not provide any security or authentication. It is solely the mechanism to manage peers in different networks to act as they were in one. If security is needed (as the case normally is), PE-to-PE tunnels can be IPsec tunnels.

7 Conclusion

The biggest challenge in BGP/MPLS VPN implementation seems to be the security. The solution is scalable, addressing is working and it does not affect other functions (at least significantly)[25]. The lack of encryption and authentication is a severe weakness since nowadays customers want to build their VPN over insecure Internet. On the other hand, it is possible to make BGP/MPLS secure to the same level as traditional Layer 2 implementations (ATM and Frame Relay)[26]. The problems of BGP/MPLS VPN are the mis-configurations of the system and attacks inside the "core". All security analysis assumes the system is configured correctly and the core is secure.

Layer 2 implementations traditionally offer good security and Layer 3 implementations scalability. Layer 2 implementations offer also more control over the encryption and authenticity since there are no underlying services reading the packets. Also if non-IP packets need to be transported, Layer 2 is the choice. But as it is stated in [15], there is no reason to say why VPN would not work in BGP when considering the features it is promised to provide.

As said earlier, the implementation of VPWS in BGP is still in the early phase and it is not just to say whether BGP is applicable place or not for it. Adding per-wire attributes for communication between to endpoints will increase the overall complexity of BGP and therefore making it less suitable for VPWS. I believe we will see more per-wire attributes. The world is not ideal and some modifications have to be made always for the theory to meet real world challenges.

BGP seems to fit to some extent for VPLS as also stated in [15]. One assertion states that VPLS is a routing protocol and the existing path selection mechanism can be used as is.

The counter-assertion is that VPLS is not a routing protocol since the path of the data coming to a PE will depend on the route to the PE, and the route is determined by ordinary IP routing.

The evolution of BGP has been in recent days clearly towards more general transport mechanism. The reasons are obvious. BGP is in wide use and has proven to be good solution for routing. It can be extended with moderate effort to do other functions than routing.

In my opinion, BGP is right place to implement an application if the application is connected to routing and BGP can be used to handle the logic, such as in case of VPN. Of course one might argue that every application has something to do with routing. The distinction has to be made in the basis of explicit use of routing in the application. In VPN, (extended) BGP can be used directly to handle the problem of overlapping namespace.

BGP should not be considered as a general application protocol. In my opinion, the use BGP should be carefully considered if it is wanted to do something else than routing. What is the limit of extending a routing protocol? There is no black or white, true or false when the question is whether BGP should be used in some application. The overall complexity of the protocol increases every time new application is added. It is tempting to use BGP as a mean of distributing applications but without strict control, BGP will become a complete mess.

References

- [1] Y. Rekhter, T. Li A Border Gateway Protocol 4 (BGP-4) RFC 1771, IETF Network Working Group, March 1995.
- [2] D. Estrin, Y. Rekhter, S. Hotz A Unified Approach to Inter-Domain Routing RFC 1322, IETF Network Working Group, May 1992
- [3] Cisco Systems Inc. Border Gateway Protocol E-Book, Cisco Systems Inc., 18 December 2003 [cited 22 March 2004] Available from World Wide Web <<http://tinyurl.com/3d5ub>>
- [4] P. Traina Autonomous System Confederations for BGP RFC 1965, IETF Network Working Group, June 1996
- [5] R. Chandra, P. Traina BGP Communities Attribute RFC 1997, IETF Network Working Group, August 1996
- [6] S. Hares Evolution of Inter-Domain Routing presentation slide, Monday June 23 2003 [cited 22 March 2004] Available from World Wide Web: <<http://www.cair.org/internetworking03/am12.htm>>
- [7] T. Bates, R. Chandra, D. Katz, Y. Rekhter Multiprotocol Extensions for BGP-4 RFC 2858, IETF Network Working Group, June 2000
- [8] S. Deering, R. Hinden Internet Protocol, Version 6 (IPv6) Specification RFC 2460, IETF Network Working Group, December 1998

- [9] Y. Rekhter, T. Li, S. Hares A Border Gateway Protocol 4 (BGP-4) INTERNET-DRAFT, IETF Network Working Group, November 2003 [cited X February 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-ietf-idr-bgp4-23.txt>>
- [10] A. Heffernan Protection of BGP Sessions via the TCP MD5 Signature Option RFC 2385, IETF Network Working Group, August 1998
- [11] J. G. Scudder, C. Appanna Multisession BGP INTERNET-DRAFT, IETF Network Working Group, November 2003 [cited 22 March 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-scudder-bgp-multisession-00.txt>>
- [12] S. R. Sangli, D. Tappan, Y. Rekhter BGP Extended Communities Attribute INTERNET-DRAFT, IETF Network Working Group, August 2003 [cited 22 March 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-ietf-idr-bgp-ext-communities-06.txt>>
- [13] A. Lange Flexible BGP Communities INTERNET-DRAFT, IETF, March 2004 [cited 22 March 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-lange-flexible-bgp-communities-02.txt>>
- [14] S. Agarwal, U.C. Berkeley, Timothy G. Griffin BGP Proxy Community Community INTERNET-DRAFT, IETF, January 2004 [cited 22 March 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-agarwal-bgp-proxy-community-00.txt>>
- [15] D. Meyer (editor) Operational Concerns and Considerations for Routing Protocol Design .. Risk, Interference, and Fit (RIFT) INTERNET-DRAFT, IETF Network Working Group, February 2004 [cited X February 2004] Available from World Wide Web: <<http://ietfreport.isoc.org/ids/draft-ietf-grow-rift-01.txt>>
- [16] VPN Consortium VPN Technologies: Definitions and Requirements white paper, VPN Consortium, January 2004 [cited 22 March 2004] Available from World Wide Web: <<http://www.vpnc.org/vpn-technologies.html>>
- [17] S. Kent Security Architecture for the Internet Protocol RFC 2401, IETF Network Working Group, November 1998
- [18] B. Patel, B. Aboba, W. Dixon, G. Zorn, S. Booth Securing L2TP using IPsec RFC 3193, IETF Network Working Group, November 2001
- [19] Tima Mangan Using Frame Relay for a VPN International Journal of Network Management, John Wiley & Sons Ltd, 2001 [cited 22 March 2004] Available from World Wide Web: <<http://tinyurl.com/2dcen>>
- [20] E. Rosen, Y. Rekhter BGP/MPLS VPNs RFC 2547, IETF Network Working Group, March 1999.
- [21] E. Rosen, A. Viswanathan, R. Callon Multiprotocol Label Switching Architecture RFC 3031, IETF Network Working Group, January 2001
- [22] L. Andersson, E. C. Rosen (Editors) L2VPN Framework INTERNET DRAFT, IETF Network Working Group, October 2003 [cited 22 March 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-ietf-l2vpn-l2-framework-03.txt>>
- [23] K. Kompella (Editor), Y. Rekhter (Editor) Virtual Private LAN Service INTERNET-DRAFT, IETF Network Working Group, January 2004 [cited X February 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-ietf-l2vpn-vpls-bgp-01.txt>>
- [24] K. Kompella (Editor) Layer 2 VPNs Over Tunnels INTERNET-DRAFT, IETF Network Working Group, January 2004 [cited 22 March 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-kompella-l2vpn-l2vpn-00.txt>>
- [25] Eric C. Rosen Applicability Statement for BGP/MPLS IP VPNs INTERNET-DRAFT, IETF Network Working Group, October 2003 [cited X February 2004] Available from World Wide Web: <<http://www.ietf.org/internet-drafts/draft-ietf-l3vpn-as2547-03.txt>>
- [26] M. Behringer Analysis of the Security of BGP/MPLS IP VPNs INTERNET DRAFT, IETF Network Working Group, 27 January 2004 [cited 22 March 2004] Available from World Wide Web <<http://www.ietf.org/internet-drafts/draft-behringer-mpls-security-06.txt>>