Border Gateway Protocols

Sadeta Krijestorac,	Marc Beck,	Jonathan Bagby
Morehead State University	University of Louisville	Florida Atlantic University
s.krijestor@moreheadstate.edu	marcbeck1982@yahoo.com	bagby@fau.edu

Abstract

A Border Gateway Protocol is a path vector routing protocol that coordinates the routing of packets through multiple administrative domains by computing routes between every IP address the packet passes. Certain routers, called BGP speakers, are assigned to run the protocol. BGP speakers across different Autonomous Systems (AS) are interconnected in order to exchange routing information. BGP supports a feature called multihoming, which means connecting to multiple ISPs from different routers or points in the network. However, BGPs still have several serious security vulnerabilities, which are currently being addressed. We discuss Pros and Cons of BGP and possible security enhancements.

Introduction

The Border Gateway Protocol (BGP) can be seen as the core interdomain routing protocol of the Internet. It is an inter-autonomous system routing protocol designed for TCP/IP networks which maintains a table of IP network prefixes that designate network reachability among autonomous systems. BGP is a path vector protocol which makes routing decisions based on paths and network policies instead of using conventional Interior Gateway Protocol (IGP) metrics. The main role of a BGP system is to exchange network reachability information with other BGP systems. In this paper we provide an overview of how BGP works, its purpose, and how it interacts with other components of the Internet as well as advantages and disadvantages of BGP over alternative protocols.

Overview of operation

The Internet is a very large-scale decentralized network consisting of smaller networks. When a packet is sent across the Internet it may pass through multiple networking administrative domains, so-called Autonomous Systems (AS). The interdomain routing of all AS's on the Internet is coordinated by the Border Gateway Protocol (BGP) running on routers that connect the AS's. The task of BGP is to compute routes between every AS and every IP address that a packet is passing on its way from one computer to another [1]. BGP is the interdomain routing protocol used to exchange reachability information between AS's on the Internet. To choose best routes, BGP allows each AS to override distance based metrics with policy based metrics [1].

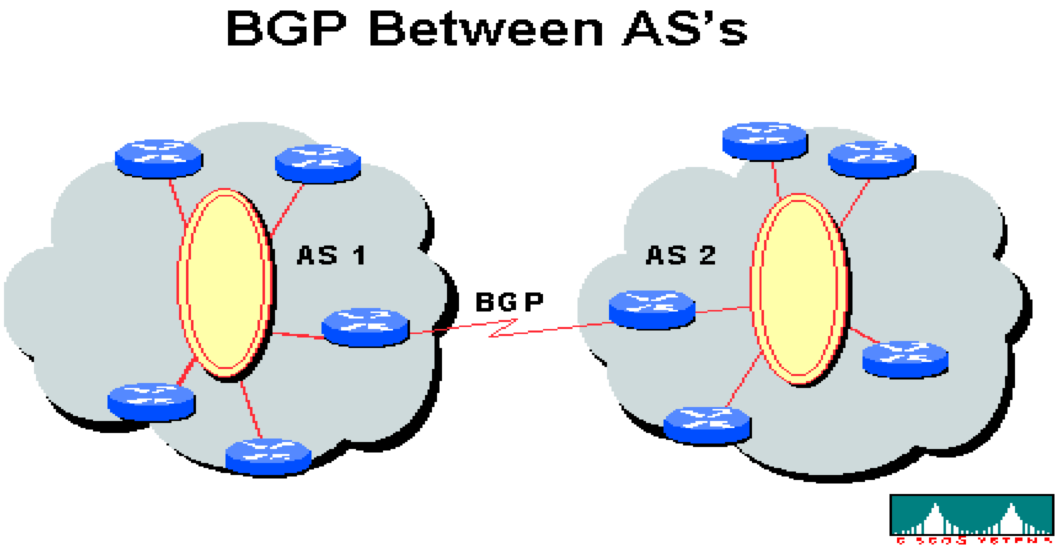

Figure 1: Chart of BGP Between AS's

Model and terminology

The Internet Engineering Task Force (IETF) created BGP as RFC 1771 and service providers first introduced it in the early 1990s as a scalable, standardized scheme to route traffic between the AS's of their customers and other service providers [2]. In order to create a BGP network, certain routers need to be assigned to run the protocol. Because they speak the BGP "language", these routers are referred to as BGP speakers. To actually create the BGP internetwork, the BGP speakers bordering each AS are physically connected to one or more BGP speakers in other AS's, ignoring any topological differences. The direct connection between them permits them to exchange information about the AS's to which they belong. BGP speakers are most often connected to multiple other speakers, which provide more direct paths to different networks for better efficiency. This also offers redundancy, allowing the Internet to deal with either device or connection failures. It is likely for a BGP speaker to have neighbor relationships with other BGP speakers both within and outside its own AS [2].

Initialization of routes

In network-layer reachability information (NLRI) aggregation, routing data to a given network in a given AS is passed along by BGP speakers in a chain fashion. Each BGP speaker in the chain appends information about its own identity and the preceding AS in the chain. As the AS routing data passes through the Internet, augmented by the list of AS's that have been passed so far, BGP forms an AS path to prevent routing loops. Once the desired topology has been defined, network administrators can determine the optimal paths and begin to set policies establishing which network destinations and communities of network destinations can exchange information [3].

Properties of the protocol

BGP is a path vector routing protocol. Each route description has several components, such as the list of prefixes being withdrawn or added, the AS path to be followed in reaching the prefix, and the address of the next router along the path [4]. The initial data-flow across a BGP backbone fills the complete BGP routing table and it gets updated incrementally when the routing tables of the other routers change. A BGP speaker must retain the current versions of all of its peers' BGP routing tables for the duration of the connection, because BGP does not refresh the entire BGP routing table and only updates changes instead. Routers periodically send keepalive messages to verify that connections are still working. BGP nodes communicate via the Transmission Control Protocol (TCP). BGP guarantees that networks within an external AS are reachable before exchanging any information by using a combination of internal BGP peering among the AS's routers and by redistributing BGP routing information to its interior gateway protocols [4].

Performance evaluation

For cost or performance reasons, it is often necessary for AS's to control the flow of their interdomain traffic. The technique of AS-Path prepending is actually useful to point out that a backup link should best be avoided if possible, but it is not easy to use it for balancing incoming traffic. AS-Path prepending is used for multihoming, which means connecting to multiple ISPs from different routers or points in the network. Quoitin, Pelsser, Bonaventure, and Uhlig have used large-scale simulations to evaluate the BGP decision process and AS-Path prepending in the Internet [5]. They found out in their simulations that the tie-break rules of the BGP decision process account for the selection of 30-50% of the routes in the global Internet. In order to control the flow of incoming packets accurately, an AS needs to be able to predict which route a distant AS will select. This prediction is very difficult to make, because the AS's knowledge of the entire Internet topology and the routing policies is often insufficient. Even if the complete topology was known, predicting the outcome of the tie-break rules of the BGP decision process would still be very complicated. Based on this analysis, the current BGP-based techniques seem not to be appropriate to control the incoming packet flow. It is suggested that changes to the Internet architecture might be necessary to achieve this kind of control [5].

Performance measure

Packet delivery is the most important performance measure for routing protocols, since this is the primary purpose of routing. The hop count can also be used as performance measure for BGP to determine the end-to-end path.
The path with the fewest links between a source and a destination will be chosen.
An ideal routing protocol should adapt rapidly to any change in topology and deliver packets as long as any path to the destination is available. Zhang et al. examined the packet delivery performance in a network running the BGP routing protocol when a destination may be disconnected from time to time [6]. Existing BGP proposals to improve convergence could negatively impact packet delivery during transient failures [6]. Most currently available routing protocols usually take seconds, or even up to several minutes, for converging after a failure. In that time, some packets may already be on their way to their destinations and new packets might have been sent. These packets can encounter routing loops, delays, and losses. There is currently not much information available about how many of them actually arrive at their destination and how many get lost during routing convergence periods [7].

Pros'

One of the greatest advantages of BGP is that corporate users can set up flexible connections between their corporate network and multiple Internet Service Providers (ISPs). For example, enterprise users can multihome and they can also set up BGP routers to automatically reroute traffic among two or more ISPs for load-sharing or backup purposes.

Two major features distinguish BGP from other routing protocols:

- It uses aggregation as a way of disseminating NLRI across routers.
- It uses path attributes for implementing routing policies [8].

Cons'

BGP has been found to be vulnerable to attacks and misconfigurations [9]. The cause of this problem is that BGP depends on information to update routing tables that is difficult to verify. Corrupted routers can add false information to the messages they transmit which other routers then use and further propagate when uncorrupted routers send extensions of these forged messages. It is easy to imagine how many serious security problems a successful compromise of a router can cause throughout the Internet [9].

Optimal Applications (topology, architecture, Layer 1 Medium)

BGP is able to connect any internetwork of AS's no matter what topology these systems use. It can handle any possible topology (full mesh, partial mesh, chain, etc) as well as changes to the topology that may occur over time when systems connect or disconnect.

The only requirement is that at least one router in each AS is able to run BGP and that this router is connected to at least one other AS's BGP router. BGP is completely unaware about what happens within the AS because it is autonomous. This means each AS has its own internal topology and set of routing protocols that it uses to make its own decisions to determine routes. BGP takes only the data that it receives from an AS and shares it with other AS's.

Alternative Protocol

Several solutions have been suggested by numerous researchers to address BGP's severe security issues. One such alternative is Secure BGP (S-BGP) which uses DSA to provide route authentication. S-BGP's actual deployment is still being prevented by several performance issues such as processing latencies and space problems like increased message size and memory cost. Zhao et al. designed aggregated path authentication schemes by combining two efficient cryptographic techniques: signature amortization and aggregate signatures [9]. They proposed constructions for aggregated path authentication that substantially improve efficiency of S-BGP's path authentication on both speed and space criteria. Their performance evaluation shows that the new schemes are efficient enough to overcome the space obstacles and offer a realistic and practical solution for BGP's- security [9].

Conclusion

BGP has been an integral part of the Internet architecture for almost two decades now. It has evolved since then in order to adapt to changes in technology, performance requirements, and security concerns. A great amount of effort has been undertaken to add new features to the original specifications. That shows that even new additions can be made and existing problems can be solved. Despite proposals for finding a replacement it seems that BGP will most likely evolve further in years to come in order to meet the demands of its users.

References

[1] Sharad Agarwal, Chen-Nee Chuah, Supratik Bhattacharyya, and Cristophe Diot, “The Impact of BGP Dynamics on Intra-Domain Traffic”, ACM SIGMETRICS Performance Evaluation Review Volume 32, Issue1 June 2004 http://research.microsoft.com/~sagarwal/sigmetrics04.pdf.

[2] Griffin, Timothy G. and Wilfong, Gordon, "An Analysis of BGP Convergence Properties", Computer Communication Review, a publication of ACM SIGCOMM, volume 29, number 4, October 1999, pp. 277--288,

[3] Meiyuan Zhao, Sean W. Smith, and David M. Nicol, "Aggregated Path Authentication for Efficient BGP Security" Dartmouth Computer Science Technical Report TR2005-541, May 2005. http://www.ists.dartmouth.edu/library/175.pdf.

[4] Jintae Kim, Steven Y. Ko, David M. Nicol, Xenofontas A. Dimitropoulos, and George F. Riley “A BGP Attack against Traffic Engineering”, Simulation Conference, 2004. Proceedings of the 2004 Winter, Volume 1, Issue 5, Dec. 2004, pp. 318 – 326.

[5] Bruno Quoitin, Cristel Pelsser, Olivier Bonaventure, and Steve Uhlig “A performance evaluation of BGP-based traffic engineering”, Intl. Journal of Network Management, Volume 15, Issue 3, May 2005, pp. 177 – 191.

[6] Beichuan Zhang, Vamsi Kambhampati, Mohit Lad, Daniel Massey, and Lixia Zhang, ”Identifying BGP Routing Table Transfers”, Conference Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement , 2002, pp 243 http://www.cs.arizona.edu/~bzhang/paper/05-minenet-mct.pdf.

[7] Atif Khan “Border Patrol--BGP”, Packet Magazine Archives, 1998, http://www.cisco.com/warp/public/784/packet/oct98/6.html

[8] Charles M. Kozierok 2005 TCP/IP guide, No Starch Press.

[9] Dan Pei, Lan Wang, Daniel Massey, S. Felix Wu and Lixia Zhang, “A study of packet delivery performance during routing convergence”, Dependable Systems and Networks, 2003. Proceedings. of IEEE International Conference on Dependable Systems and Networks , 22-25 June 2003, pp183 – 192.

Biography

SADETA KRIJESTORAC is currently Assistant Professor at the Department of Applied Engineering and Technology at Morehead State University, KY. Dr. Krijestorac has over 15 years of experience as a software development engineer and educator. She is president of the Kentucky Academy of Science, Engineering Division.

MARC BECK is currently a PhD student at Speed School of Engineering of the University of Lousville, KY. He holds a Bachelor degree and Master degree from the Department of Applied Engineering and Technology at Morehead State University.

JONATHAN BAGBY is currently associate Professor at the Department of Electrical Engineering at Florida Atlantic University, FL. Dr. Bagby has over 20 years of experience as an engineer, developer, and educator. His area of expertise is electromagnetic computation and fiber and computer communication and networking.