On BGP Communities

Benoit Donnet, Olivier Bonaventure
Universite catholique de Louvain – CSE Department
Belgium
{firstname.name}@uclouvain.be
This article is an editorial note submitted to CCR. It has NOT been peer reviewed. Authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

This paper focuses on BGP communities, a particular BGP attribute that has not yet been extensively studied by the research community. It allows an operator to group destinations in a single entity to which the same routing decisions might be applied. In this paper, we show that the usage of this attribute has increased and that it also contributes to routing table growth. In addition, we propose a taxonomy of BGP community attributes to allow operators to better document their communities. We further manually collect information on BGP communities and tag it according to our taxonomy. We show that a large proportion of the BGP communities are used for traffic engineering purposes.

Categories and Subject Descriptors

C.2.2 [**Network Protocols**]: Routing Protocols

General Terms

Taxonomy, Measurement

Keywords

BGP, community attribute, traffic engineering

1. INTRODUCTION

The *Border Gateway Protocol* (BGP) [1] is the current inter-domain routing protocol. A domain or *Autonomous System* (AS) is either a single network or a group of networks that is under the control of a single administrative entity, typically an Internet Service Provider (ISP) or a large organization with independent connections to several other ASes. BGP is used to exchange routing information between ISPs. Routes learned via BGP contain *attributes* that are notably used to determine the best route to a prefix when multiple paths exist to reach this prefix. Typical BGP attributes are LOCAL_PREF, MED (suggestion to an external AS regarding the preferred route into the AS that is advertising the metric), or AS_PATH (list of ASes that the route advertisement has traversed). In this paper, we focus on a particular optional BGP attribute defined in RFC 1997: the *BGP communities attribute* [2].

The BGP communities attribute provides a way of grouping destinations into a single entity, named *community*, to which similar routing decisions might be applied. A BGP communities attribute is composed of one or more 32 bits numbers. These numbers are structured as follows: the high-order 16 bits represent an AS number, while the low-order 16 bits define the semantic of the value. Each AS can use the 2^{16} communities whose high-order 16 bits are equal to its own AS number.

The BGP communities attribute has found several usages. For instance, it might be applied in multi-homing routing, as defined in RFC 1998 [3]. By replacing the AS-based customization of the BGP LOCAL_PREF attribute, it allows a client to influence the LOCAL_PREF attached to its routers in its provider's network. This can be used to indicate that routes advertised over a (primary) high-speed link should be preferred over routes advertised over a (backup) low-speed link. Standard values for BGP communities are described in RFC 1997 [2] and their applications are discussed by Foster [4].

The contributions of this paper are twofold. Based on Routeviews [5] and RIPE [6] BGP routing tables, we first evaluate the importance of the BGP communities attribute. We show that, since September 2004, there is an increasing usage of the BGP communities attribute. We further explain that this situation has an impact on the BGP routing tables growth due to the memory consumed by BGP communities. Next, starting from Quoitin and Bonaventure's work [7], we propose a taxonomy of the BGP communities attribute that is based on three categories: *inbound* (description of a route's entry point), *outbound* (traffic engineering purpose), and *blackhole* (for security). We collect publicly available BGP community information and build a database tagged using our taxonomy. Using this database, we show that the majority of the BGP communities found in RIPE and Routeviews routing tables are used for traffic engineering purposes.

The remainder of this paper is organized as follows: Sec. 2 analyzes the importance of BGP communities in routing tables; Sec. 3 discusses our BGP communities taxonomy; Sec. 4 evaluates how our taxonomy can be used to tag BGP communities in BGP table dumps; Finally, Sec. 5 concludes and discusses further directions.

2. IMPORTANCE OF BGP COMMUNITIES IN ROUTING TABLES

In this section, we evaluate the importance of the BGP communities attribute in routing tables. We base our evaluation on publicly available routing table dumps from the Routeviews [5] and RIPE [6] projects. For this study, we only take into account IPv4 routes.

Our Routeviews dataset is composed of four BGP routers: Dixie (Japan), Equinix (United States of America), ISC (United

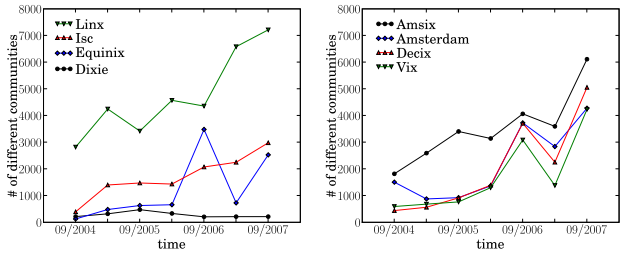

(a) Routeviews (b) RIPE

Figure 1: Evolution of BGP communities over time

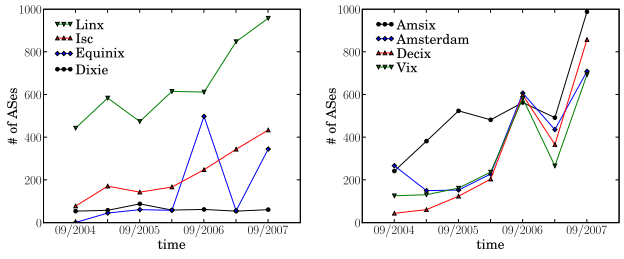

(a) Routeviews (b) RIPE

Figure 2: Evolution of ASes using BGP communities over time

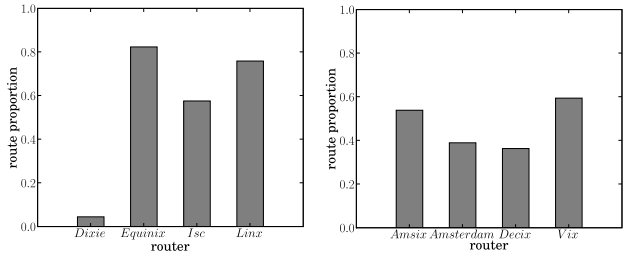

(a) Routeviews (b) RIPE

Figure 3: Proportion of routes carrying BGP communities (Sept. 2007)

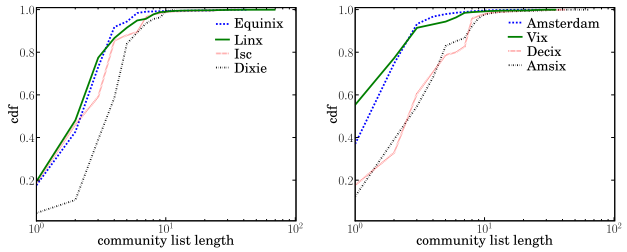

(a) Routeviews (b) RIPE

Figure 4: BGP community list length in BGP tables (Sept. 2007)

States of America), and Linx (United Kingdom). We consider four routers in the RIPE dataset: AMSIX (The Netherlands), Amsterdam (The Netherlands), DECIX (Germany), and VIX (Austria). The largest BGP table is Linx with more than 5,000,000 routes stored. On the opposite, the smallest one is Dixie, with roughly 695,000 routes.

We first look at the evolution of the BGP communities attribute usage. Starting from September 2004 to September 2007, we consider BGP table dumps every six months for each router in both datasets. Fig. 1 shows the evolution of the different BGP communities values stored in routing tables (i.e., the number of communities attributes stored) over time.

Except for the Dixie router (Routeviews dataset – see Fig. 1(a)), we can see that there is an increasing usage of BGP communities over time. For instance, the AMSIX router (RIPE dataset) encounters nearly three times more different BGP communities values in September 2007 than in September 2004. In addition, we notice a drop in the growth of BGP communities usage in March 2007 for the RIPE dataset.

If the BGP communities usage has increased over time, it would be interesting to know if this growth is caused by a few ASes or if the number of ASes making use of BGP communities also increases. This is the purpose of Fig. 2. If we look at both Fig. 1 and 2 in parallel, we notice that there is a growth in the number of ASes using BGP communities and this growth follows the same trend as the evolution of BGP communities values considered. This means that using BGP communities is becoming more and more popular among ASes.

Let us now have a more detailed look at the current situation and focus on a recent BGP table dump (i.e., last dump of September 2007). Fig. 3 shows the proportion of routes that carry at least one BGP communities value. A value of one means that all the routes stored in the BGP table contains this attribute. On the other hand, a value of zero indicates that no route contains the communities attribute. Note that, on the contrary to Fig 1, a single communities attribute may contain several 32 bits values (this is known as *BGP communities list*).

On average, one can see that about half of the routes contain BGP communities. We notice, however, two extreme cases: for Equinix (see Fig. 3(a)), more than 80% of the routes contain BGP communities while for Dixie (see Fig. 3(a)), only around 5% of the routes make use of BGP communities.

If more or less half of the routes carry a BGP communities attribute, it would be interesting to know the contribution of this attribute to the memory consumption growth into routing tables.

Fig. 4 shows the cumulative distribution of the length of the communities values list (i.e., the number of 32 bits values contained in each BGP communities attribute), in log-scale, for the last dump of September 2007. If we look at the Routeviews dataset (Fig. 4(a)), we see that more than 50% of the routes carrying BGP communities contain more than two 32 bits values. We further notice that a small proportion of routes carrying BGP communities information in the Linx router contain more than 60 different 32 bits values. The situation is somewhat identical for the RIPE dataset (Fig. 4(b)) with the difference that, for two routers (Amsterdam and VIX), only 25% of the routes containing BGP communities information carry more than two 32 bits values.

If we now have an idea of the number of BGP communities carried in a route, we have to evaluate the memory they consumed on BGP routers. To achieve this, we have to consider how a BGP implementation stores BGP communities. To this aim, we worked as follows: we sorted all received BGP communities received so that for each communities list, the communities items are in increasing order. We next calculated a 32 bits hash on this list and maintained

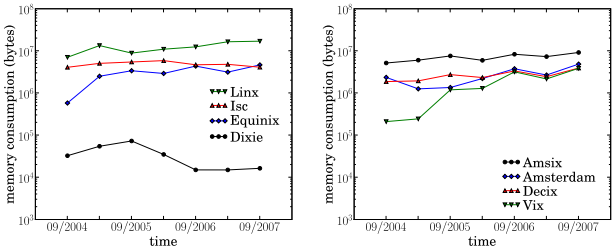

(a) Routeviews (b) RIPE

Figure 5: BGP communities memory usage in routing tables

a counter indicating the number of different hash value obtained. Based on that, the memory size will be the set of BGP communities list plus a 32 bits pointer for each prefix. This models the behavior of the quagga BGP daemon [8].

Fig. 5 shows the evolution of the BGP communities memory consumption over time. The horizontal axis gives the time while the vertical axis, in log-scale, shows the memory consumption expressed in bytes.

Except for the Dixie router (see Fig. 5(a)), every router presents an increase in the memory consumption produced by the BGP communities usage. In particular, for the Vix router (see Fig. 5(b)), the BGP communities consumes, in September 2007, eighteen times more memory space than in September 2004. The usage of BGP communities is thus also a factor impacting the routing table size growth.

3. BGP COMMUNITY TAXONOMY

In Sec. 2, we have seen the importance of BGP communities in routing tables. We believe it is important to provide a taxonomy of BGP communities to better describe their usages. This taxonomy and the associated BGP communities documentation is publicly available[1].

Fig. 6 shows our classification of BGP communities, an extension of Quoitin and Bonaventure's work [7]. This classification mainly consists of three categories: the *inbound*, the *outbound*, and the *blackhole* communities, each one being divided into various subtypes. The inbound communities are discussed in Sec. 3.1, outbound communities in Sec. 3.2, and blackhole communities in Sec. 3.3.

3.1 Inbound Communities

Inbound communities refer to communities added or used when a route is received by a router on an eBGP session. We divide the inbound communities into two categories, as depicted in Fig. 6: *LocalPref* and *Routes tagging*. The first type of communities is used to set the LOCAL_PREF attribute in the AS receiving the route. An example from AS174: 174:10 sets the LOCAL_PREF to 10. This corresponds to the BGP communities usage described in RFC 1998 [3]. Note that BGP communities that can tweak a route's preference should be carefully engineered and consistently implemented in order to avoid unintended routing, as explained in RFC4264 [9]. The second type of communities is used by an AS to indicate the location where the route was received from an external peer.

A route might be tagged in different ways. In our taxonomy, we consider the four most common types of tag:

1. *IXP.* This type of communities indicates the interconnection point where a route was learned. For instance, 4589:14901, defined by AS4589, indicates that the route was received by AS4589 at Decix.

2. *Type of Peer.* ASes have agreements between each others, depending on the type of relationship they maintain such as peer, customer, provider, or transit. An AS might define a community value for describing this relationship. This simplifies the configuration of filters on BGP routers. For instance, the community 286:900, from AS286, tags transit routes.

3. *Geographic.* An AS might need to indicate the geographic location where the route was received. This location might be quite general (i.e., a continent), or more precise (i.e., a city). For instance, 86:4013, defined by AS86, informs that the route was learned by AS86 in Brussels.

4. *AS.* This community indicates the AS from which a route was learned. This is different from the Type of Peer type as it does not specify the relationship between the ASes. For instance, AS15997 tags all routes learned from TeleGlobe with 15997:1080.

Most inbound communities values in use today have been defined independently by different operators. Recently, Dave Meyer [10] proposed recommended encodings for most of these values. However, our analysis tends to indicate that these encodings are not currently used.

3.2 Outbound Communities

Outbound communities are used by a router to filter BGP announcements for traffic engineering purposes. A community is inserted by the originator of the route in order to influence its redistribution by downstream routers. The outbound communities affect thus the *redistribution of routes*. We divide the routes redistribution communities into two categories, as shown on Fig. 6: *Announcement* and AS_PATH *prepending*.

Regarding the announcement, the community is attached to a route to indicate whether the route should or should not be announced to a specified peer or at a specified interconnection point. For instance, the community 8938:4000, defined by AS8938, means that AS8938 routers should not announce to any upstream provider.

AS_PATH prepending aims at prepending n times to the AS_PATH when announcing the route to a specified peer. For example, values 286:1n, defined by AS286, prepends n (where $1 \leq n \leq 4$) times AS286 to European peers. This prepending action is performed by a router belonging to AS286.

3.3 Blackhole

The blackhole category refers to a particular BGP community used by an ISP to block packets. One can distinguish three types of blackholing by their scope. With the *global* blackholing, packets are blocked everywhere in the ISP networks. With the *border* blackholing, packets are blocked at the ISP border routers. Finally, with the *upstream* blackhole, packets are blocked by the ISP's transit provider before they enter its own network [11, 12].

[1]See `http://inl.info.ucl.ac.be/communities`.

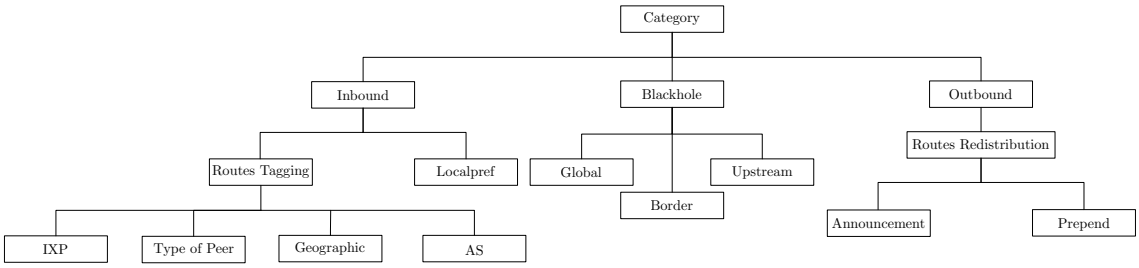

Figure 6: BGP Communities taxonomy

Category	Category	Type	Proportion (%)	Total (%)
Blackhole			0.22%	0.22%
Inbound	Localpref		1.62%	26.25%
	Routes tagging	Interconnection Point	1.82%	
		Type of Peer	6.33%	
		Geographic Location	9.11%	
		AS	6.08%	
	Unknown		1.26%	
Outbound	Routes redistribution	Announcement	25.9%	73.52%
		AS_PATH prepending	47.43%	
	Unknown		0.16%	

Table 1: Repartition of the collected BGP community documentation

Blackhole communities have been introduced to face security issues. In particular, it provides a way to block denial-of-service (DoS) attacks.

These communities are used only inside ISPs and should not be distributed on the global Internet.

4. DOCUMENTING BGP COMMUNITIES

Based on the taxonomy described in Sec. 3, we built a database that contains description of BGP communities. Unfortunately, there is no standard and formal way of documenting the BGP communities values used by ISPs. Some provide information on their web site, others include these definitions in the comments of their RPSL objects stored on whois servers. We used AS numbers available in the RIPE and Routeviews datasets from September 2007 and manually collected available BGP communities information from whois servers and ISPs websites for these ASes. This collection was performed during Fall 2007. Our database is freely available to the networking community.[2] It currently contains 27,180 entries spread into the various categories as described in Table 1.

Looking at Table 1, we see that the main usage of BGP communities is related to traffic engineering (73%), i.e., outbound communities, and, more particularly, to AS_PATH prepending. This means that traffic engineering techniques are more complex than typically thought. It also means that AS-level topology models, such those recently introduced by Mulbauer et al. [13, 14], by Gao [15], or Mahadevan et al. [16], should probably be revised in the light of the BGP communities attribute usage.

Note that for inbound and outbound communities, a tiny proportion of BGP documentation was not precise enough to decide to which category it belongs. We label such a community as "unknown".

We then used this database as input for classifying BGP communities found in BGP table dumps from September 2007. The results of this classification (i.e., the proportion of classified and unclassified BGP communities per router) are shown on Table 2. We see that we are able to classify, on average, 22% of the BGP communities, for both datasets. We are currently discussing with operators to collect more information on the BGP communities that they have defined.

Fig. 7 runs into the details of BGP communities we were

[2]See `http://inl.info.ucl.ac.be/communities`.

Dataset	Router	Classified	Unclassified
Routeviews	Dixie	18.31%	81.69%
	Equinix	26.90%	73.10%
	Linx	23.50%	76.50%
	ISC	18.49%	81.51%
	Average	**22.80%**	**78.20%**
RIPE	AMSIX	21.30%	78.70%
	Amsterdam	24.95%	75.05%
	DECIX	21.90%	78.10%
	VIX	24.04%	75.96%
	Average	**23.05%**	**76.95%**

Table 2: Classified vs. unclassified (Sept. 2007)

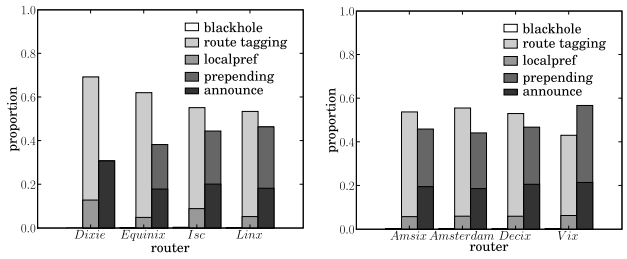

(a) Routeviews (b) RIPE

Figure 7: Repartition of classified BGP communities (Sept. 2007)

able to classify. Most of the recognized communities are inbound communities (55% on average) and, in particular, those that describe geographic location. Regarding the outbound communities, most of them are AS_PATH prepending. Finally, a very tiny proportion of recognized communities belongs to the blackhole category.

5. CONCLUSION

In this paper, we focused on the BGP communities attribute. This attribute provides a way of grouping destinations into a single entity, named *community*, to which similar routing decisions might be applied.

In this paper, we have seen that the BGP communities attribute is more and more used by operators. This motivated us for providing a more structured way of describing BGP communities values. We proposed a taxonomy based on three main categories: inbound, outbound, and blackhole communities. We explained that the inbound communities refer to a description of the route's entry point, while outbound communities follow mainly traffic engineering purposes, the blackhole category mainly focusing on security issues,

We manually collected more than 27,000 BGP communities attribute information and tagged them according to our taxonomy. We showed that a large proportion (around 75%) of the communities are outbound communities. It means that current efforts in providing a model for the AS-level topology must necessarily consider the BGP communities attribute.

When applying our database on recent BGP table dumps, we were able to tag, on average, 22% of the communities. Most of them were inbound communities.

Future work might reveal how we can use this BGP communities taxonomy and database for additional usage. For instance, any change in inbound communities might reveal a topology change and be considered as a trigger-event for launching a traceroute-like exploration.

We also believe that it would be interesting to have a more precise way of describing BGP communities instead of the informal documentation currently found on web sites or RPSL. The taxonomy we proposed in this paper is a first step towards this end. We are currently working on a C-BGP model [17] that could be used by operators to describe their networks as well as to detail the specificities of their BGP communities.

Acknowledgements

This work was supported by the European-founded 034819 OneLab project. Authors would like to thank the NANOG community as well as RIPE for helping them in collecting BGP communities documentation. Authors are indebted to the reviewer for his proof-reading of the paper and Bruno Quoitin for our fruitful discussion on BGP communities.

6. REFERENCES

[1] Y. Rekhter and T. J. Watson, "A border gateway protocol 4 (BGP-4)," Internet Engineering Task Force, RFC 1771, March 1995.

[2] R. Chandra and P. Traina, "BGP communities attribute," Internet Engineering Task Force, RFC 1997, August 1996.

[3] E. Chen and T. Bates, "An application of the BGP community attribute in multi-home routing," Internet Engineering Task Force, RFC 1998, August 1996.

[4] K. Foster, "Application of BGP communities," *Cisco Internet Protocol Journal (IPJ)*, vol. 6, no. 2, pp. 2–9, September 2003.

[5] University of Oregon, "Route views, University of Oregon Route Views project," see http://www.routeviews.org/.

[6] RIPE NCC, "Routing information service (RIS)," October 1999, see http://www.ripe.net/np/ris/.

[7] B. Quoitin and O. Bonaventure, "A survey of the utilization of the BGP community attribute," Internet Engineering Task Force, Internet Draft (work in progress) draft-quoitin-bgp-comm-survey-00.txt, February 2002.

[8] "Quagga routing suite," 1999, see http://www.quagga.net.

[9] T. Griffin and G. Huston, "BGP wedgies," Internet Engineering Task Force, RFC 4264, November 2005.

[10] D. Meyer, "BGP communities for data collection," Internet Engineering Task Force, RFC 4384, February 2006.

[11] B. R. Green and D. McPherson, "ISP security: Deploying and using sinkholes," June 2003, NANOG 28, see http://www.nanog.org/mtg-0306/sink.html.

[12] J. Soricelli and W. Gustavus, "Options for blackhole and discard routing," October 2004, NANOG 32, see http://www.nanog.org/mtg-0410/soricelli.html.

[13] W. Mulbauer, S. Uhlig, B. Fu, M. Meulle, and O. Maennel, "In search of an appropriate granularity to model routing policies," in *Proc. ACM SIGCOMM*, August 2007.

[14] W. Mulbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an AS-topology model that captures route diversity," in *Proc. ACM SIGCOMM*, September 2006.

[15] L. Gao, "On inferring autonomous system relationships in the internet," in *Proc. IEEE Global Internet Symposium*, Nov. 2000.

[16] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, k. claffy, and A. Vahdat, "The Internet AS-level topology: Three data sources and one definitive metric," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 17–26, January 2006.

[17] B. Quoitin and S. Uhlig, "Modeling the routing of an autonomous system with C-BGP," *IEEE Network*, vol. 19, no. 6, pp. 12–19, November 2005.