Protocol Overhead versus Router Processing and Memory Tradeoffs in Evaluating BGP Inbound Soft Reset Mechanisms for Tactical IP Networks

Isil Sebuktekin and John Haluska
Telcordia Technologies, Inc
Piscataway, NJ
isil@research.telcordia.com and jhaluska@telcordia.com

Pete Moyer
Pollere Inc,
Menlo Park, CA
pete@pollere.net

Kevin Adams
Lockheed Martin IS&GS
San Jose, CA
kevin.adams@lmco.com

***Abstract*—Border Gateway Protocol (BGP) is the inter-domain routing protocol of choice across the Global Information Grid (GIG) as in the commercial Internet. There is a future need and motivation to extend BGP to connect the large-scale future military networks of the Army, Navy and Air Force to the GIG as independent Autonomous Systems (ASes) using satellite and wide area networking (WAN) technologies. These networks are expected to be highly mobile and multi-homed to the GIG. Higher BGP protocol activity and policy use, both in the inter- and intra-domain, is anticipated for such mobile ASes compared to the largely static commercial counterparts. Wireless bandwidth being a scarcer and more volatile network resource, emphasizes the need to manage protocol overheads among routers in dynamic mobile network environments. This is in contrast to the primary concern of optimizing router packet processing and memory utilization in significantly higher speed provider networks of the Internet. In consideration of the different dynamics expected, we analyze BGP "soft resets" as one small but high-impact operational property of BGP for the tactical Internet, specifically the "inbound soft resets" and the different methods for implementing them. We support our comparative analysis with empirical results, also evaluating impacts on BGP Route Reflectors with multiple peers.**

I. INTRODUCTION

Border Gateway Protocol version 4 (BGP-4) [1] is the ubiquitous inter-domain routing protocol globally used in the commercial and tactical Internet. There is increased momentum in the development of the Global Information Grid (GIG) to extend the footprint of BGP to connect large-scale stationary and mobile military networks of the Navy, Army and Air Force as independent Autonomous Systems (ASes) administratively managed by their respective organizations.

Tier 1 and Tier 2 military networks of the GIG are expected to operate largely over wireless wide-area network infra-structures and to be multi-homed to the GIG to achieve the reliability and redundant connectivity needs and requirements of tactical ASes. Geographically distributed large-scale networks such as Army's WIN-T and Navy's ADNS would likely be extensively multi-homed across multiple ASes to better handle mobility and wireless impediments.

Military networks would also likely be using typical BGP scalability and security solutions within and at ingress to their ASes such as Route Reflectors [2], route-flap damping [3], ingress filtering [4][5], outbound route filtering capability [6][7], etc. despite the difficulties introduced by mobility and multi-homing. It is important to understand the impacts of internal and external BGP (iBGP[2] and eBGP[3]) dynamics within and in between large-scale mobile tactical networks with different connectivity characteristics and policy needs than the wireline networks of the commercial Internet. In this paper, we analyze a high-impact but small aspect of the BGP protocol operations for such tactical networks, namely the BGP "soft resets" [8] required to en-act any potential changes in BGP policy configurations.

While BGP "soft resets" are not inherently disruptive to data plane traffic forwarding, they are management actions, which result in re-advertisement of specific sets of routes impacted by the new BGP configuration or policy. Consequently, bursts of subsequent control plane, i.e., routing protocol overheads may adversely affect the share of

[1] Prepared through collaborative participation in the Transformational Satellite Communications System (TSAT) Mission Operations System (TMOS) Program, sponsored by the U.S. Air Force, Contract FA8808-06-C-0003. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

[2] BGP connections between internal BGP peers, i.e., BGP routers belonging to the same AS

[3] BGP connections between external BGP peers, i.e., BGP routers belonging to different ASes

Digital Object Identifier: 10.4108/ICST.COLLABORATECOM2009.8373
http://dx.doi.org/10.4108/ICST.COLLABORATECOM2009.8373

bandwidth available for data plane traffic for a period until convergence is reached. Bandwidth capacity is typically the scarcest resource in tactical IP networks in contrast to wireline networks. With this consideration in mind, we analyze in this paper how some overhead savings can be realized with particular uses of soft resets, more specifically in the use of "BGP inbound soft resets" between BGP peers.

II. Overview of BGP Soft Resets

A BGP session must be reset in order to enact any policy changes on a router that may potentially be configured during an active BGP session. A BGP policy change may be as simple as an update in a BGP path attribute or a metric configured on a router. It is only after a BGP reset that the new policy may start influencing the routing decisions BGP peers make to reach the routes advertised by the other.

BGP resets could be accomplished by means of either soft or hard resets [8], although hard resets are disruptive and undesirable as they result in complete session tear-down and the subsequent establishment of an independent BGP session. They are hardly used in operational environments unless extreme actions such as a reboot are needed. BGP "soft resets" are an old protocol refinement over BGP "hard resets", whereby the updated routes are exchanged, but the BGP session is not re-initialized as with a hard reset, facilitating non-disruptive routing policy updates. A soft reset could be instantiated either inbound or outbound, depending on the type of the BGP policy change and which set of inbound or outbound routes it is going to affect:

- A BGP "outbound soft reset" on a router causes it to advertise its updated BGP routes to its peer, enabling the new BGP policies configured on the router to apply such as a new BGP metric or path attribute associated with the BGP routes.
- A BGP "inbound soft reset" on a router on the other hand, results in application of updated policy on the router to those routes advertised by the peer. These may be the routes already-received from the peer and stored in memory in anticipation of dynamic changes in policy, or an inbound soft reset may potentially trigger a request to the peer to re-advertise its routes so that the updated BGP policy can now be applied.

There are two possible ways of performing BGP "inbound soft resets" with differing overhead impacts on a peer-to-peer BGP session. The choice of the method for use is simply controlled by subtle variations in router configurations:

1. BGP Route Refresh [9]: With this method an "inbound soft reset" on a router causes it to send a BGP ROUTE REFRESH (type 5) packet to its peer. The BGP peer in turn, responds by re-advertising all or a specified subset of its advertised-routes to the requesting router. Subsequently, the local router applies the new policy to the recently received routes from the peer. To use this method, both BGP peers must be pre-configured to negotiate BGP session capabilities during the initial BGP establishment phase and the Route Refresh capability must be supported by each peer. This is often the de-fault configuration in typical COTS routers supported by the capabilities negotiation in BGP OPEN process.
2. BGP Soft Reconfiguration [9]: This is the fall-back and older soft reset method when Route Refresh capability is not supported by a BGP peer. Soft Reconfiguration requires the local router to keep at all times a copy of the peer's advertised-routes, enabling any updates in local inbound routing policy to be directly applied without needing a Route Refresh and subsequent re-advertisement of the peer's routes to be transmitted through the network.

Please note that BGP outbound soft resets used to enact an outbound routing policy changes do not have similar multiple variations of methods as applicable to inbound soft re-sets. The advertisement of updated routes is automatically triggered towards the peer upon enactment of the policy change by the outbound soft reset on the local BGP router.

Between the BGP Route Refresh and Soft Reconfiguration as inbound soft reset mechanisms, there is naturally a trade off in the resultant BGP routing overhead on the network with the BGP Route Refresh extension [9] versus the additional memory required on a BGP router to store the peer's current version of advertised BGP routes throughout the lifetime of the connection [1].

Intuitively, Soft Reconfiguration based inbound soft reset provides routing protocol overhead savings over the Route Refresh method, which we will compare empirically further below. Reduced protocol overhead is clearly a more critical consideration over resource-constrained wireless links subject to variable network impairments, typically packet losses that slows TCP over which BGP is transported.

Memory for routing table storage is also an important consideration, but is often less critical than conserving wireless bandwidth resources. This is much different than the primary resource conservation considerations in high-speed wireline networks. In fact the BGP Route Refresh method was introduced later over the older method of soft recon-figuration for the purpose of conserving memory in back-bone routers in service provider networks when memory historically used to be a much more expensive resource. Even though the memory costs have been much reduced since then, Route Refresh has typically been the default method for inbound soft resets on COTS routers.

Multiplicative memory storage impacts may need careful consideration however when the BGP router to store all its BGP neighbors' routing tables has a large number of peers, or is a Route Reflector (RR) with a large number of RR Client peers. The Route Reflectors used in an AS's iBGP configurations for scaling intra-domain routing, typically have multiple iBGP RR Client peers constituted by all or a subset of the border routers of the AS. We compare below the two methods of implementing inbound soft resets using empirical results, also discussing the potential multiplicative affects of routing overhead and router memory.

Digital Object Identifier: 10.4108/ICST.COLLABORATECOM2009.8373
http://dx.doi.org/10.4108/ICST.COLLABORATECOM2009.8373

III. An Empirical Setup To Compare BGP Inbound Soft Reset Methods

To compare the impacts of different methods of BGP in-bound soft resets, we utilize empirical measurements over a small emulated multi-AS GIG testbed, exercising BGP multi-homing relationships between the ASes. The portion of the testbed we use to discuss the differences in inbound soft reset mechanisms is shown in Figure 1. It highlights a Transit Service Provider AS (TRANSIT AS) and two neighboring ASes (GIG AS and USER AS), where one of them (USER AS) is multi-homed to the Transit Service Provider AS (TRANSIT AS) at multiple locations, two of which used in this empirical test are shown in Figure 1.

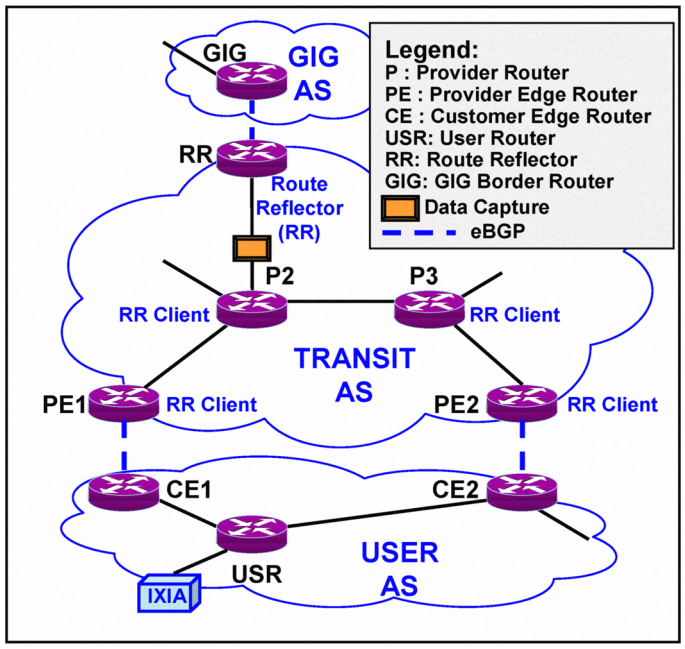

Figure 1. A Multi-AS Testbed Emulating the GIG

The TRANSIT AS in this testbed has a number of Provider Routers, or P Routers in the network core, and a number of Provider Edge, or PE Routers bordering the other two neighboring ASes. In our particular experiments to compare the two mechanisms, one PE Router labeled as the RR borders the GIG AS in Figure 1, and is configured as an iBGP Route Reflector (RR) for the other PE routers of the TRANSIT AS. Both the P and PE routers in this AS are RR Client iBGP peers of the RR and they all border the other neighboring AS, i.e., the USER AS, multi-homed to the TRANSIT AS. Each of these PE routers provides a different routing path to reach the users in or behind the USER AS for any user in or behind the TRANSIT AS, and the GIG AS that the TRANSIT AS serves. Likewise there are different paths in the reverse direction to reach any GIG or TRANSIT AS user from the USER AS.

To demonstrate the impact of different inbound soft reset methods, we experiment with a single iBGP path attribute, the local-preference (or LOCAL_PREF) metric in this example to control which of these two different paths (via the two distinct PE-CE pair of routers in Figure 1) is preferred by the RR, relative to the other path for all inter-domain routing from the GIG AS to the USER AS. For simplicity, we apply the same LOCAL_PREF routing policy on RR to reach all the prefixes in or behind the USER AS, although in practice different policies can be configured for different sets of prefixes.

Approximately 3,370 routes (emulated by an IXIA generator shown in Figure 1, as well as the local subnet prefixes and loopbacks in the USER AS) are advertised via eBGP in this testbed by each of the USER AS CE routers to their PE router eBGP peers in the TRANSIT AS. These routes are subsequently advertised via iBGP by the RR client PE Routers to the RR, and then via eBGP to the GIG AS, as reachable inter-domain routes.

IV. Summary of Experiments That Require an Inbound Soft Reset

The relative values of the LOCAL_PREF metric configured on the RR for each RR client peer determine the preferred routing path from the RR towards the USER AS routes, in other words which RR Client iBGP peer is being selected as the egress router. This is simply achieved by configuring an inbound route-map on the RR for each RR client PE peer to set the "local-pref" path attribute value as desired for the PE-advertised prefixes which then influence the BGP best path routing decisions taken at the RR to reach USER AS routes via the preferred PE egress.

The highest LOCAL_PREF value towards a PE peer determines the most preferred route to go through the corresponding PE router. For this experiment, we initially configured the highest LOCAL_PREF value on RR towards PE1. This controls the preferred routing path from the GIG AS to the USER AS to go through the routers RR, P2, and PE1 in the TRANSIT AS and CE1 in the USER AS. We set a smaller LOCAL_PREF to be tagged onto PE2-advertised prefixes as part of this initial RR configuration.

In our experiment to compare the impacts of different inbound soft reset mechanisms, the initial BGP path preference policy configured on the RR is illustrated in Figure 2, along with the preferred path towards the USER AS routes.

A network event or a change in usage policy may necessitate switching the preferred routing path from across PE1 to across another PE to reach all or a subset of USER AS prefixes. By updating the relative values of the LOCAL_PREF metrics on the RR per iBGP peer, it is possible to update the initial iBGP path preference policy on the RR to now use a different path via a different PE Router, to reach the same USER AS routes.

In this example, let's assume that the routing path across PE1 has become no longer prefer-able and that the network administrator reconfigures the LOCAL_PREF metric for PE1 to a value less than the value configured for PE2. This updated BGP path preference policy configured on RR is shown in Figure 3, along with the new preferred path towards USER AS routes.

Digital Object Identifier: 10.4108/ICST.COLLABORATECOM2009.8373
http://dx.doi.org/10.4108/ICST.COLLABORATECOM2009.8373

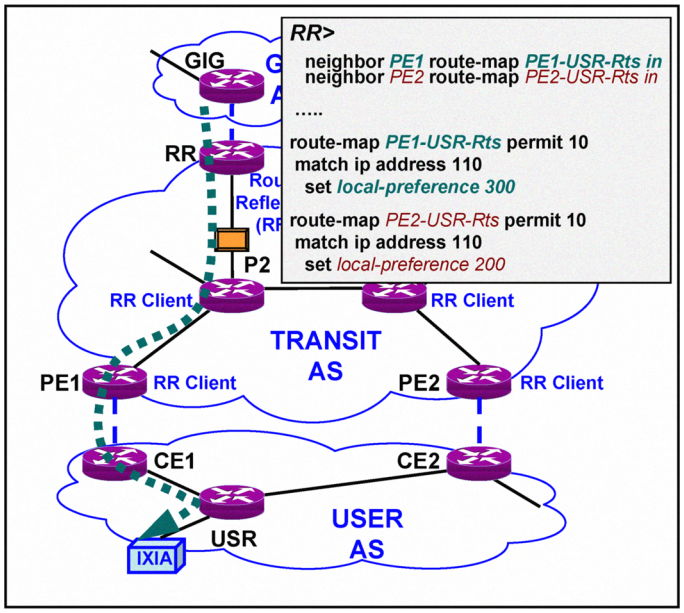

Figure 2. Initial Path Preference Policy Configuration

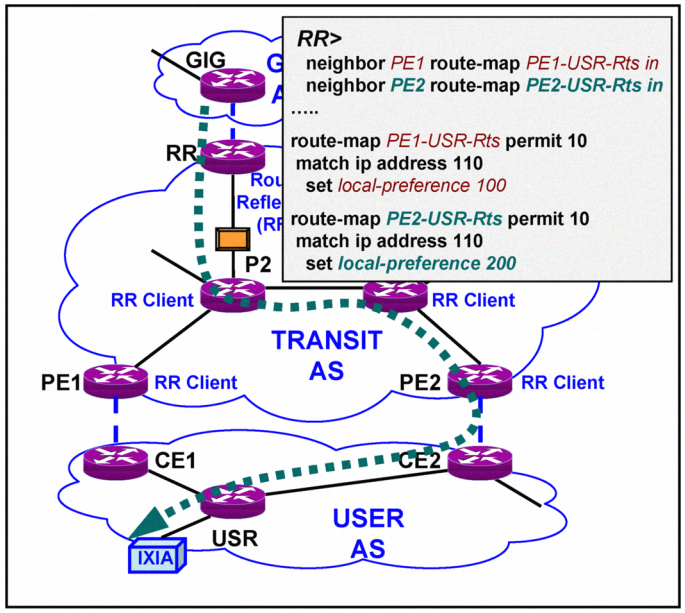

Figure 3. Updated Path Preference Policy Configuration

Configuration statements in Figure 2 and Figure 3 provide as examples, the configurations used in this experiment. In order for this new preferred path policy to take effect, the updated configuration on the RR should be succeeded by an inbound soft reset on the RR towards the affected PE peer (PE1 for which the LOCAL-PREF value has just been changed). The inbound soft reset causes enacting of the change in the value of the LOCAL_PREF metric and results in a series of BGP UPDATE messages and protocol processing until routing converges to the new best path.

The amount and order of BGP UPDATE messages exchanged as a result of the inbound soft reset varies as discussed below, depending on the type of the BGP inbound soft reset mechanism configured and used between peers.

V. Summary of BGP Overhead Results Comparing Inbound Soft Reset Methods

With the Route Refresh method[4], an inbound soft reset on the RR triggers a BGP ROUTE-REFRESH message to be sent from RR to PE1, resulting in PE1 to re-advertise to RR all the USER AS prefixes as shown in Figure 4 below. RR then applies the new LOCAL_PREF value to PE1 advertized routes and informs all its RR client iBGP peers of the new best route towards USER AS via PE2. The values in parentheses in Figure 4 below represent the actual number of BGP bytes[5] sent.

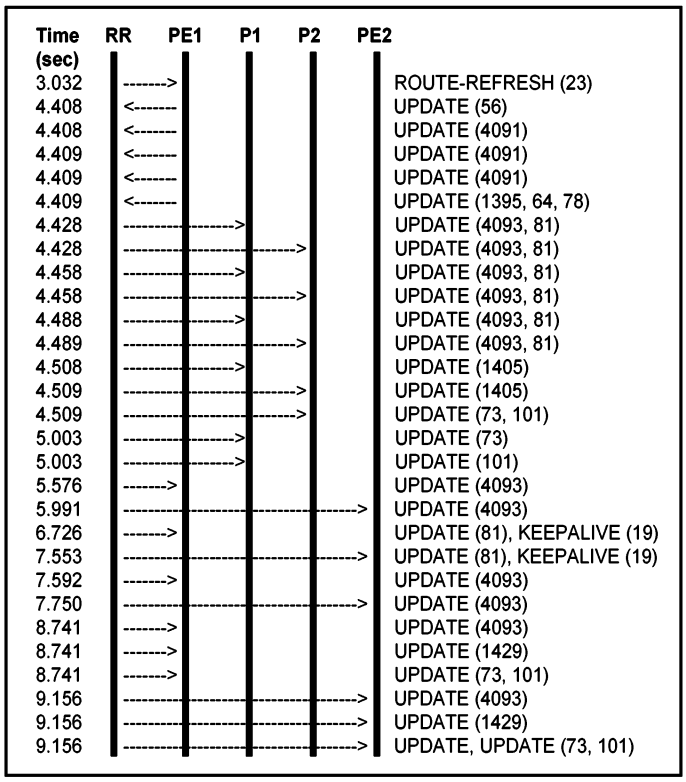

Figure 4. BGP Flow Graph Following the LOCAL_PREF Policy Update and Route Refresh based BGP Soft Reset

[4] Route Refresh Capability must be negotiated by the peers during initial connection establishment (BGP OPEN) phase.

[5] For those readers interested in details of BGP UPDATE messages and different BGP protocol attributes [1], multiple byte values listed in parenthesis for some UPDATE messages in Figure 4 correspond to groups of prefixes with a different BGP ORIGIN attribute in a single BGP UPDATE message (e.g., IGP, INCOMPLETE, etc.). For example, the first such occurrence (1395, 64, 78) in Figure 4 is for a PE1-to-RR UPDATE message with 1395 Bytes of prefixes corresponding to IXIA-advertized routes showing INCOMPLETE ORIGIN, 64 Bytes of prefixes listing peer CE1's local subnets showing IGP as their ORIGIN and 78 Bytes of other USER AS router subnets also with IGP ORIGIN. Last message on the other hand in Figure 4 assembles two small, but separate UPDATE messages in a single TCP segment, each with a different BGP ORIGINATOR_ID attribute besides different BGP ORIGIN attributes.

Digital Object Identifier: 10.4108/ICST.COLLABORATECOM2009.8373
http://dx.doi.org/10.4108/ICST.COLLABORATECOM2009.8373

With the BGP soft reconfiguration method[6] on the other hand, no ROUTE-REFRESH messages are exchanged and no BGP UPDATE messages are sent from PE1 to RR as illustrated in Figure 5, corresponding to nearly 14,000 bytes of control plane overhead savings (13,889 bytes to be exact) for the 3,370 IPv4 prefixes advertised. For the total 65,881 BGP Bytes sent with the Route Refresh method, this represents a 21% savings[7] in overhead in this particular example of 3,370 prefixes. RR simply tags the locally stored USER AS routes that it has previously received from PE1 with the new LOCAL_PREF value. Subsequently, BGP best path selection algorithm favors PE2 with higher LOCAL_PREF value. As in Figure 4, the values in parentheses in Figure 5 likewise represent the number of bytes in the BGP messages.

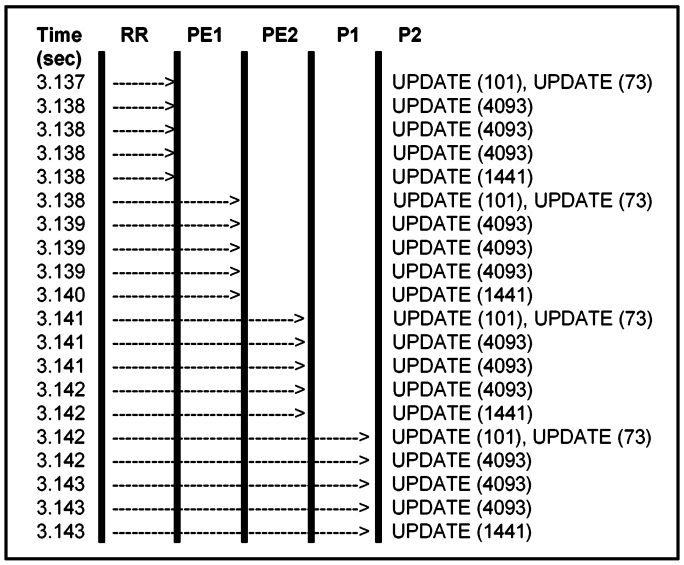

Figure 5. BGP Flow Graph Following the LOCAL_PREF Policy Update and Soft Reconfiguration based Reset

In either case, the updated LOCAL_PREF values causing PE2 to be the preferred next hop router for the RR to reach the 3,770 USER AS routes, triggers BGP advertisements of the updated best path routes from the RR to each RR client peer, whereby within each method the same amount of BGP bytes are sent to each RR client P and PE router peer in advertising the USER AS prefixes.

As illustrated within the parentheses in Figure 4 and Figure 5, each BGP UPDATE message displayed in a single line in these flow graphs represents a BGP message of up to 4093 bytes in length. Each of these large BGP messages corresponds to three to four reassembled TCP segments, each carried in a 1514-byte Ethernet frame across our testbed links and routers configured with TCP Path MTU Discovery[8] in order to avoid TCP segmentation across different domains. Back-to-back small BGP UPDATE messages (e.g., those advertising prefixes with different path attributes) and BGP KEEPALIVE messages are also assembled and sent together within the same TCP segment, and subsequently in the same Ethernet frame, when appropriate as shown in some lines in Figure 4 and Figure 5.

One could easily observe in the Figure 5 flow graph that the RR sends all BGP UPDATES corresponding to 3,370 network prefixes to its RR client peers in a strictly sequential manner when soft reconfiguration is used. Further exploration into the packet contents also show that the USER AS prefixes transmitted within the Network Layer Routing Information (NLRI) field of the BGP UPDATE messages are sorted from the largest prefix to the smallest in the case of the soft reconfiguration method. In contrast, the prefixes within the NLRI field are somewhat arbitrarily ordered in the case of the Route Refresh method as is the order of the initially advertised prefixes by the IXIA generator, and subsequently by PE1 to RR, in response to the ROUTE-REFRESH request from RR. This may also potentially be implementation dependent however; our experiments were on a single-vendor testbed.

Consequently, the soft reconfiguration method appears to provide better efficiency in assembling and transmitting short BGP packets together, resulting in overhead savings compared to Route Refresh based inbound soft resets (5 versus 6 BGP messages sent to each PE client peer). The Route Refresh based method as a result seem to roughly introduce an additional 100-200 Byte BGP control plane overhead per RR client peer in RR's subsequent reflection of the routing updates for the 3,370 USER AS routes.

VI. Comparing CPU / Memory Utilization Between Inbound Soft Reset Methods

RR being the router on which the inbound soft reset is performed and as the router responsible for reflecting updated routing information to multiple RR client peers, it is interesting to compare how the RR performance fairs in terms of its CPU and memory utilization between the two methods we analyze in this paper following the inbound soft reset on the RR. Before we discuss its CPU and memory utilization performance specifically, we first compare the packet transmission times shown in Figure 4 and Figure 5.

The 1.5-second time delay observed between the ROUTE-REFRESH message sent to PE1 and the BGP UPDATES received by PE1 in Figure 5 is due to round trip delay configured in our testbed and should not be included in this comparison. We therefore compare only the packet transmission times from the RR. These timestamps and packet sizes shown in Figure 4 and Figure 5 are obtained from packet captures between RR and P2 in our testbed, recorded by means of the Wireshark [10] network protocol analyzer tool.

With Soft Reconfiguration based inbound soft resets, the transmission of BGP updates from RR onto the link towards

[6] To implement soft reconfiguration based inbound soft resets, the configurations in Figure 2 and Figure3 above need to be supplemented with two additional statements, one per RR client PE peer. These are "neighbor [PE1, PE2] soft-reconfiguration inbound" statements. It is important to make sure that the default Route Refresh capability does not override this which may require explicitly turning it off on some routers.

[7] Note that overhead savings would vary by the number of prefixes that need to be advertised and may be more or less than the 21% savings observed in this example.

[8] BGP runs over TCP utilizing its reliable transmission mechanisms.

Digital Object Identifier: 10.4108/ICST.COLLABORATECOM2009.8373
http://dx.doi.org/10.4108/ICST.COLLABORATECOM2009.8373

RR client peers is back-to-back and faster (about 6 msec with Soft Reconfiguration versus over 4.5 seconds with Route Refresh) than with Route Refresh based resets when they are already sorted in memory.

With Route Refresh based inbound soft resets, the transmission of BGP Update messages is interleaved across the RR client peers and spread over time as they are received from PE1. This is not a distinguishing performance characteristic between the two methods however when the CPU utilization performance on the RR is actually compared between the two methods of BGP inbound soft resets upon LOCAL-PREF policy update as presented in Figure 6 and Figure 7.

CPU utilization on RR is shown in Figure 6 with Route Refresh based reset and in Figure 7with Soft Reconfiguration based reset, used to enact the LOCAL-PREF policy update on RR. As observed in these two figures, CPU utilization peaks in each case between 9th and 12th seconds after the reset, and appears busier in the case of Soft Reconfiguration based BGP inbound soft reset.

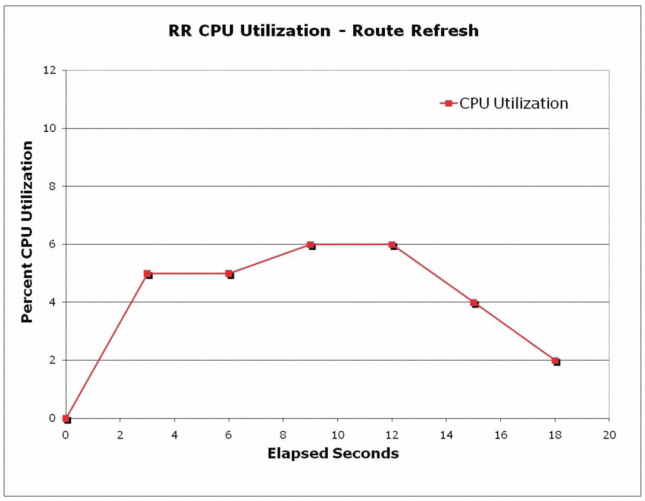

Figure 6. RR's CPU Utilization upon LOCAL_PREF Policy Update and Route Refresh based Inbound Soft Reset

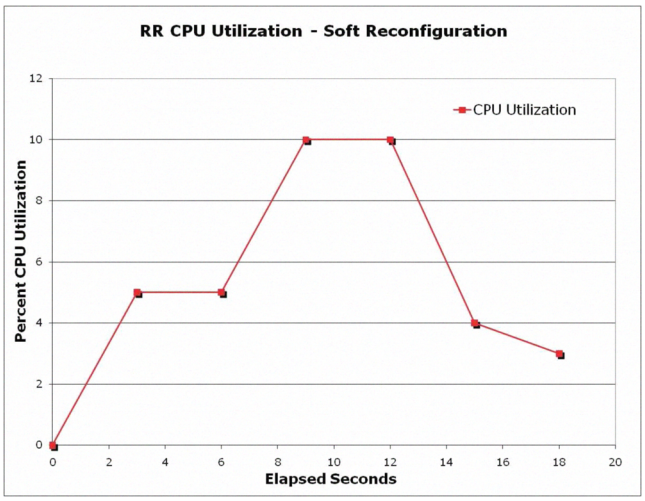

Figure 7. RR's CPU Utilization upon LOCAL_PREF Policy Update and Soft Reconfiguration based Reset

The BGP memory and router memory processing on RR is also higher with soft reconfiguration based BGP inbound soft reset method as shown in Figure 8 and Figure 9 below. Higher memory utilization[9] is naturally expected since RR needs to keep a local copy of all routes advertized by all its RR client peers during the lifetime of these BGP peering connections and fetches them off of memory upon soft reconfiguration based inbound soft resets. The difference is most observable on the memory utilization of BGP router processes, but is negligibly small with only few RR Client peers and few thousands of routes advertised from each peer. This may have to be considered more carefully with larger numbers of peers and per-peer advertized routes.

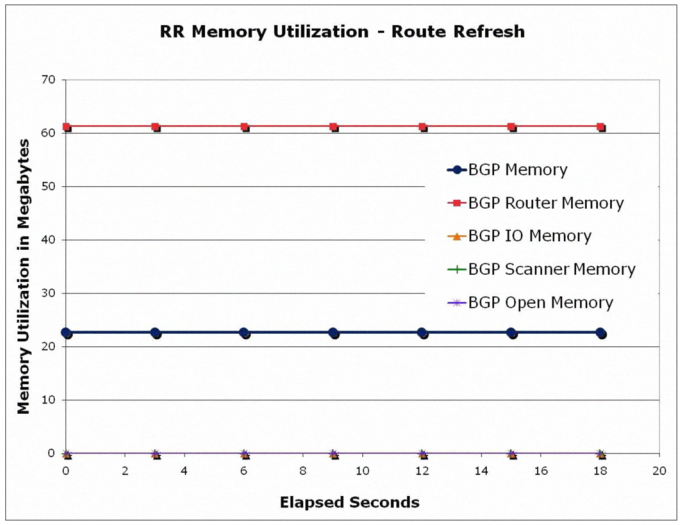

Figure 8. RR's Memory Utilization upon LOCAL_PREF Policy Update and Route Refresh based Soft Reset

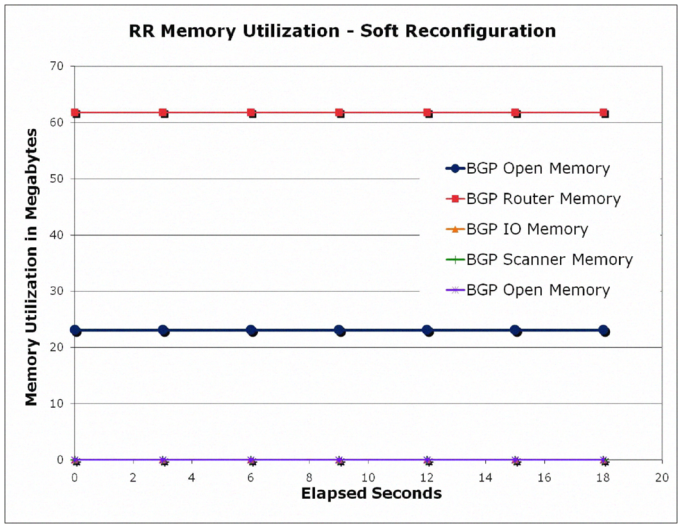

Figure 9. RR's Memory Utilization upon LOCAL_PREF Policy Update and Soft Reconfiguration based Soft Reset

[9] The memory utilization graphs presented actually represent slightly higher numbers of RR client iBGP peers and few more thousands of advertised routes, which have been purposely left out in this paper to simplify the experiments we discuss to a bare minimum.

Digital Object Identifier: 10.4108/ICST.COLLABORATECOM2009.8373
http://dx.doi.org/10.4108/ICST.COLLABORATECOM2009.8373

Table 1 below provides the actual memory utilization values measured in Megabytes for the various BGP processes for the two soft reset methods analyzed and computes the difference in memory utilization in Megabytes between the Soft Reconfiguration method and the Route Refresh method for this particular exercise involving 3,770 prefixes.

Table 1. Memory Utilization in Megabytes for BGP Processes

Memory Utilization in Megabytes for Various BGP Processes					
Soft Reset Method Used	**BGP Memory**	**BGP Router Memory**	**BGP IO Memory**	**BGP Scanner Memory**	**BGP Open Memory**
Route Refresh	22,670,953	61,290,360	7,296	10,296	7,228
Soft Reconfiguration	(max) 23,105,933	61,815,064	47,468	10,296	0
	(min) 23,105,685				
Difference in Memory Utilization	(max) 434,980	524,704	40,172	0	-7,228
	(min) 434,732				

Note that memory utilization numbers would change by the number of prefixes involved in BGP routing exchanges, and are only provided here for informative reasons as otherwise the granularity of memory utilization data gets lost in the large scale of graphs illustrated. In short, Table 1 also confirms that Soft Reconfiguration method requires more memory than the Route Refresh method. It should be obvious to the reader that the Route Refresh method would only use more memory in the BGP Open process, caused by ROUTE-REFRESH messages.

VII. CONCLUSIONS

In this paper we discuss and empirically compare two alternative mechanisms of performing BGP inbound soft resets specifically for large-scale military networks of the GIG, namely BGP Route Refresh and BGP Soft Reconfiguration, using as example an iBGP routing path policy change involving LOCAL-PREF metric requiring an inbound BGP soft reset to enact the change.

Soft Reconfiguration was the only inbound soft reset mechanism used initially in early BGP deployments. Route Refresh mechanism was later developed and standardized in the IETF to optimize router memory processing for core Internet routers operating at optical transmission rates. As we have empirically quantified for one specific example involving inbound soft resets, the trade off between the two mechanisms is the additional BGP routing overhead on network links with Route Refresh versus the additional memory required on a BGP router with Soft Reconfiguration to permanently store its peer's advertised routes.

As our primary conclusion, we favor Soft Reconfiguration based inbound soft resets for large-scale military networks using BGP to connect to the GIG, against the conventional use preferred about a decade ago. This is mainly because the network dynamics and optimization considerations in this trade-off for tactical networks are much different than those for high-speed wire-line commercial counterparts. In particular, protocol overhead savings is more critical across resource-constrained wireless links subject to variable network impairments typically resulting in non-negligible packet losses that slow TCP over which BGP resides.

While more careful consideration is needed when very large numbers of network routes or BGP peers are involved, memory and router memory processing savings nowadays are less critical factors, given the advances and significant cost reductions in storage technology over the last decade since BGP Route Refresh mechanism was standardized.

Hence, since the Route Refresh capability has typically been the default method for inbound soft resets on COTS routers (negotiated during the initial BGP connection establishment process), there is a need to override this default configuration on BGP router peers in tactical networks. In our empirical example we also saw Soft Reconfiguration based inbound soft resets to provide additional overhead savings if the router on which the reset is performed happens to be a Route Reflector, reflecting updated BGP routes to its RR client BGP peers. This additional savings is small however, and may be router implementation specific

In either case what matters in practice is the average rate of BGP control plane traffic with respect to the overall bottleneck link bandwidth between the peers and how the forwarding of routing and data packets is handled by the intermediate routers. This is why we suggest minimizing BGP protocol overheads using Soft Reconfiguration based inbound soft resets when a reset is necessary besides as common sense, avoiding any unnecessary and frequent resets overall in tactical networks.

ACKNOWLEDGEMENTS

The authors would like to extend their special thanks to Gene Ma of the Aerospace Corporation for his review of the paper and useful technical comments that helped improve its contents.

REFERENCES

[1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, January 2006.

[2] T. Bates, E. Chen, and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)," RFC 4456, April 2006.

[3] C. Villamizar, R. Chandra, and R. Govindan, "A BGP Route Flap Damping," RFC 2439, November 1998.

[4] P. Ferguson, and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," RFC 2827, May 2000.

[5] F. Baker, and P. Savola, "Ingress Filtering for Multihomed Networks," RFC 3704, March 2004.

[6] E. Chen, and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4," RFC 5291, August 2008.

[7] E. Chen, and S. Sangli, "Address-Prefix-Based Outbound Route Filter for BGP-4," RFC 5292, August 2008.

[8] H. Berkowitz, E. Davies, S. Hares, P. Krishnaswamy, and M. Lepp, "Terminology for Benchmarking BGP," RFC 4098, June 2005.

[9] E. Chen, "Route Refresh Capability for BGP-4," RFC 2918, September 2000.

[10] Wireshark Network Protocol Analyzer, http://www.wireshark.org.

Digital Object Identifier: 10.4108/ICST.COLLABORATECOM2009.8373
http://dx.doi.org/10.4108/ICST.COLLABORATECOM2009.8373