A Configuration based Approach to Mitigating Man-in-the-Middle Attacks in Enterprise Cloud IaaS Networks running BGP

Stephen Brako Oti
Information Technology Department, Methodist University College, Ghana

Isaac Bansah
Information Technology Department, Methodist University College, Ghana

Tonny M. Adegboyega
Information Technology Department, Methodist University College, Ghana

ABSTRACT
Cloud IaaS service providers offer virtualized computing resources to enterprises over the internet. As with most internet based services, cloud service providers may need to establish BGP peering relationships with upstream/neighbor ISPs for the purposes of exchanging routing information between their respective Autonomous systems thereby making it possible for a rogue AS to carry out a Man-In-The-Middle (MITM) attack. Available literature supports the fact that BGP as an infrastructure protocol is vulnerable to MITM attacks yet a good number of proposals aimed at counteracting these attacks have not been fully implemented. Secure BGP, Secure Origin BGP and Pretty Secure BGP are all proposals which have not been fully implemented due to high overhead and invariable router load. We believe however that an existing cloud IaaS service provider could mitigate the risk of a MITM attack by optimizing their configurations and ensuring that upstream providers do a proper job filtering prefixes using a prefix-list. This paper presents a GNS-3 simulation of a MITM attack by mimicking a section of the internet and goes on to show how the application of a prefix-list can help mitigate the attack.

General Terms
Man-In-The Middle Attacks, BGP Vulnerabilities, Autonomous Systems, Cloud IaaS Services.

Keywords
BGP security, Cloud Services, Prefix-List, Secure BGP, Session Hijacking.

1. INTRODUCTION
The rapid expansion of telecommunication infrastructure alongside exponential increments in the processing and data transfer capabilities of today's personal computer precipitates the proliferation of cloud computing services around the world. Cloud computing basically refers to the ability to offer virtualized computing resources to end users or enterprises packaged as IaaS, PaaS or SaaS. Generally speaking, cloud computing from an enterprise customer's perspective is a tradeoff between cost, control and security. While the option of moving computing resources and data to the cloud offers a genuine value for money business case, there are also very legitimate considerations bordering on control over data and security that must inform the decision to move to the cloud and further shape how Cloud service providers design their networks to reduce the potential risk associated with an IaaS offering for instance.

IaaS Cloud service providers together with their enterprise customers face a multiplicity of specialized attacks targeted at either the client or the provider for every individual instance of occurrence. Enterprise customers could suffer the problem of comprised credentials and broken authentication [1] where data breaches occur as a result of weak passwords, poor key or certificate management. The huge volumes of data and other virtualized resources typically stored and provided by cloud service providers makes them an attractive target for Data breaches, Denial of Service attacks, Man-In-The-Middle attacks and a host of other vulnerabilities.

Service providers typically form BGP peering relationships with neighbor networks as part of the primary Infrastructure requirements necessary to keep the Cloud service running and reachable over the internet. This opens up a whole new chapter of vulnerabilities which could affect a cloud IaaS service due to the normal operation of the BGP4 protocol. BGP Version 4 is the default Inter-Domain routing protocol currently used to exchange routing information between autonomous systems [2]. There is documented proof of previous BGP incidents such as the YouTube incident of 2008 where Pakistan Telecom advertised a more specific route to YouTube resulting in the redirection of YouTube traffic to Pakistan Telecom causing YouTube to be unavailable for about two hours. Although this incident was not specifically intended to create a black hole or Man-In-The-Middle attack, a rogue autonomous system could set out to exploit this BGP flaw so as to bring about a more devastating attack.

To address this potential vulnerability, a number of proposals namely Secure BGP (SBGP), Pretty Secure (PsBGP) have been put forward however due to some operational and deployment issues,[3] these technologies have not been fully implemented. One of the major drawbacks with SBGP for instance is the inability of existing deployed routing infrastructure to support the additional overhead associated with the encryption SBGP works with. Another dimension of countermeasure is a proposal for a BGP prefix hijack alert system [4] which notifies prefix owners in real time when their BGP origin changes. In view of the numerous efforts and proposals put forward towards achieving greater security with BGP, we postulate that the proper use of a prefix-list in filtering routes coupled with greater care and responsibility with respect to the administration of border routers could go a long way to mitigate the MITM attack risk. It is noteworthy to mention that in the 2008 YouTube incident, PCCW Hong Kong could have helped curtail the false broadcast originating from Pakistan Telecom by filtering out the false route [3]. We demonstrate the workability of our solution with a simulation built in GNS-3 with routers running the CISCO IOS 15. This approach makes no additional demands on the existing deployed routing equipment by way of processing power and overheads as the prefix-list [12] feature can be implemented

on an existing Cisco router without necessarily upgrading the hardware.

2. RELATED WORK

In section three, we examine a number proposals put forward for securing BPG against MITM and other similar attacks pointing out the highpoints and low points.

3. CURRENT PROPOSALS FOR SECURING BGP

In developing our own approach to mitigating MITM attacks, we reviewed a number of approaches already put forward with the aim of protecting BGP from session hijacking, prefix hijacking as well as MITM attacks. Although not explicitly developed for the purpose of mitigating MITM attacks, approaches such as SBGP, PsBGP, SoBGP and Prefix Hijack Alert System (PHAS) can be adapted to protect against various BGP vulnerabilities that could affect a Cloud IaaS service provider network. In this section, we analyze the various approaches in a bid to point out their shortcomings and thereby establish the relevance and feasibility of our approach.

3.1 Secure BGP

Secure BGP addresses fundamental security issues in BGP4 by employing digital signatures for authentication, as well as the use of a PKI in validating these digital signatures [5]. The real issues being addressed here has to do with the provision of a mechanism for ISPs to validate the identity of other ASes and their ownership of specific IP Prefixes. The overall architecture of Secure BGP is anchored on three security mechanisms which are Public Key Infrastructure, a new optional BGP transitive path attribute and finally IPSEC. In the first instance, the PKI is used to provide an authentication framework necessary to validate the identity of IP Block as well as AS number owners. Furthermore, the PKI provides a mechanism to validate the identity of an AS as well as the identity of the BGP router and as to whether that router is authorized to represent the AS in the BGP peering session. In the second instance, sBGP proposes the use of attestations using the BGP Transitive Path attribute and digital signatures [6]. The attestation authorizes a nominated AS to advertise itself as the origin AS for a particular address prefix based on the use of the sBGP PKI and digital certificates used. Finally, IPSEC is suggested for securing inter-router communication paths so as to provide data and partial sequence integrity thereby making it possible for BGP routers to authenticate each other for the exchange of BGP control traffic.

Despite the very promising prospects of sBGP, issues of computational overhead associated with its implementation [7] on existing deployed BGP routers alongside the need for collaboration between several distinct bodies such as Internet Registries and Internet Service Providers have been cited as some of the reasons why sBGP has not been fully implemented.

3.2 Secure Origin BGP

Secure origin BGP (soBGP) is another effort aimed at securing BGP put forward by Russ White [8]. soBGP tries to achieve a workable balance between the computational overhead associated with the implementation of sBGP vis-à-vis the capabilities of existing deployed routing infrastructure as well as the collaborative cooperation of Internet Security Infrastructure bodies required for the success of sBGP.

The main issue of authentication; thus in the case of BGP has to do with the ability of participating entities in a BGP session to validate the identity of other ASes and also to know the kind of information they will be using to sign their data. soBGP addresses this challenge using EntityCerts which ties an AS number to a public key or set of public keys which corresponds to a private key the AS will be using to sign various other certificates. soBGP further uses an Authorization Certificate "AuthCert" to provide authorization for an AS to advertise a specific block of addresses after establishing the identity of the AS through the use of EntityCerts [8].

The avoidance of a hierarchical PKI for the validation of AuthCerts and EntityCerts is a way of simplifying the use of soBGP but could also be considered a weakness in this approach, as the derivation of authority to speak on addresses is very unclear in this model [2].

3.3 Pretty Secure BGP

psBGP is a BGP security effort put forward by Van Oorschot et al [9]. psBGP attempts to combine the best features of both sBGP and soBGP into a new proposal that provides a justifiable balance between security, performance and practicality [9]. psBGP employs two separate trust models for authenticating AS numbers as well as for the verification of properties associated with IP prefix origination. In the first instance, psBGP employs a centralized trust model or framework for the authentication of AS numbers where each AS could obtain a public key certificate from one out of a number of trusted certificate authorities essentially binding that AS number to the given public key. Binding an AS number to a specific public key is expected to provide some amount of integrity with respect to the identity of ASes considering the credibility of the Certification Authorities involved and the reduced risk of impersonation by a rogue AS. In the second instance, psBGP employs a decentralized trust model to validate the identity of IP prefix owners using a prefix assertion list which binds AS numbers to IP prefixes for a given AS and another list for the peering ASes of the given AS. psBGP appears to be needlessly complex and bears much of the characteristics of making a particular solution for the problem, rather than attempting to craft a solution within the bounds of the problem space [2].

3.4 Prefix Hijack Alert System (PHAS)

PHAS is a reactional approach proposed by Lad et al [10] aimed at creating an alert mechanism that notifies prefix owners whenever their BGP origin changes. The proposal focuses more on prefix owners finding out about potential hijacks in real time rather than preventing the hijack from taking place all together as suggested by the previous authors as in the case of sBGP, soBGP and psBGP. By providing reliable and timely notification of origin AS changes, PHAS allows prefix owners to quickly and easily detect prefix hijacking events and take prompt action to address the problem. PHAS typically monitors and analyses logs gathered from BGP collectors such as Route Views and analyses the data for changes in the BGP origin. While PHAS may succeed in notifying AS owners of potential prefix hijacks or origin changes, it does not include mechanism to validate the identity of ASes making it more reactional than preventive.

4. BGP MAN-IN-THE MIDDLE ATTACK SIMULATION

BGP does not provide protection against man-in-the-middle attacks. As BGP does not perform peer entity authentication, a man-in-the-middle attack is child's play [11]. To effectively demonstrate the feasibility of applying prefix-lists as counter

measure against BGP MITM attacks, we prepared a simulation in GNS-3 with routers running CISCO iOS 18 mimicking a section of the internet where a cloud IaaS provider is peering with BGP neighbors in the first scenario. In the subsequent scenario, we simulate a BGP MITM attack and subsequently demonstrate how the prefix-list could be applied to filter out bogus updates from a rogue AS.

4.1 Scenario 1- Cloud Service Provider Normal BGP Operation

The Target is a Cloud Based Service provider, advertising its prefix/network (20.20.20.20) to its BGP Peers. A trace from AS_100 to the Cloud Server goes through, the routers in AS_300→AS_500. This is normal operation

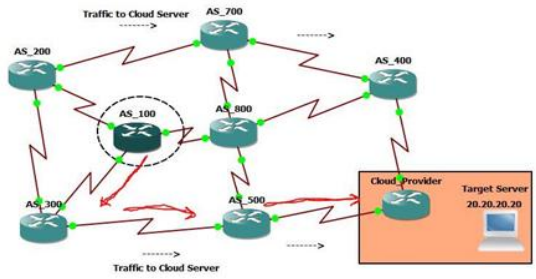

Fig 1: Shows normal traffic flow, with traffic destined for the Cloud Provider travelling to the provider's network as intended.

```
AS_100#trace 20.20.20.20

Type escape sequence to abort.
Tracing the route to 20.20.20.20

  1 41.202.1.2 0 msec 0 msec 0 msec
  2 31.202.1.2 [AS 300] 0 msec 0 msec 0 msec
  3 21.202.1.2 [AS 500] 0 msec 12 msec 0 msec
AS_100#
```

Fig 2: Shows the trace output

4.2 Scenario 2 – Malicious Routes Inserted Into BGP

The attacker now begins to inject malicious routes into BGP causing traffic destined to the Cloud Server to go through the Attacker network causing traffic destined to the Cloud Server to go through the Attacker network.

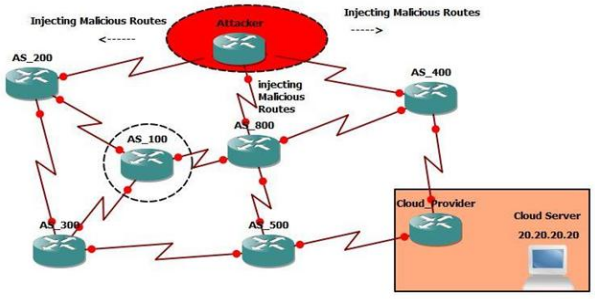

Fig 3: Shows an attacker injecting a more specific prefix into BGP and affecting the BGP table/routes on the global Internet

A trace from AS_100 now attempts to go through the Attackers network (31.202.0.1).

```
AS_100#
AS_100#trace 20.20.20.20

Type escape sequence to abort.
Tracing the route to 20.20.20.20

  1 41.202.1.2 4 msec 0 msec 4 msec
  2 31.202.1.2 [AS 300] 0 msec 4 msec 0 msec
  3 10.202.0.1 [AS 500] 4 msec 0 msec 4 msec
  4 31.202.0.1 [AS 800] 4 msec 8 msec 0 msec
  5 31.202.0.1 [AS 800] !H  !H  !H
```

Fig 4: Shows output of the trace

A BGP output from AS_100 also reveals that the path to 20.20.20.20 is going through the Attacker because he is generating a more specific prefix for that network, making its path more appealing.

```
AS_100#sh ip bgp
BGP table version is 18, local router ID is 10.10.10.10
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  0.0.0.0          41.202.0.2               0             0 200 i
*>                  41.202.1.2               0             0 300 i
*  9.202.0.0/30     41.202.0.2               0             0 200 ?
*>                  41.202.1.2               0             0 300 ?
*> 10.10.10.0/24    0.0.0.0                  0         32768 i
*  10.202.0.0/30    41.202.0.2                             0 200 300 500 ?
*>                  41.202.1.2                             0 300 500 ?
*  11.202.0.0/30    41.202.1.2                             0 300 200 ?
*>                  41.202.0.2               0             0 200 ?
*  11.202.1.0/30    41.202.0.2                             0 200 300 500 800 700 400 ?
*>                  41.202.1.2                             0 300 500 800 700 400 ?
*  20.20.0.0/19     41.202.0.2                             0 200 300 500 600 i
*>                  41.202.1.2                             0 300 500 600 i
*  20.20.20.0/22    41.202.0.2                             0 200 300 500 800 700 i
*>                  41.202.1.2                             0 300 500 800 700 i
AS_100#
```

Fig 5: Shows output of the "show ip bgp" command on AS_100

4.3 Scenario 3 – Redirecting Traffic to Avoid Black hole

Although the attacker now has traffic destined for the Cloud Server going through it, he however does not actually have that Server on his network so the traffic actually drops when it gets to him.

See pings drops from AS_100 below for the 20.20.20.20 Server.

```
AS_100#
AS_100#ping 20.20.20.20

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.20.20.20, timeout is 2 seconds:
UUUUU
Success rate is 0 percent (0/5)
```

Fig 6: Shows pings destined for the 20.20.20.20 network dropping

The Attacker fixes this by redirecting the traffic destined for the Cloud Server through AS_400. This redirection is possible because the Attacker has manipulated his BGP session with AS_400 (using AS Prepends), making the Attacker's path to the Cloud Server unappealing to AS_400. Hence the AS_400 network does not go through the Attacker network to get to the Cloud Server.

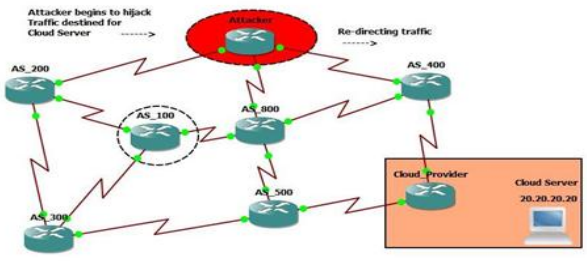

Fig 7: Shows the attacker fixing the black hole by redirecting the traffic meant for the cloud server through AS_400

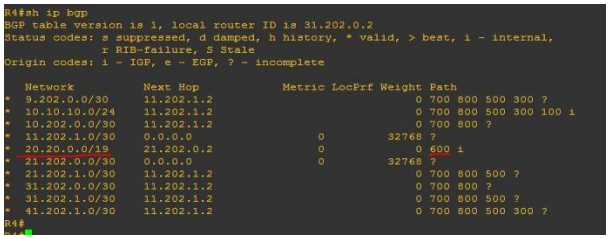

Fig 8: Shows router Output indicating that AS_400 is using AS_600 to get to Cloud Server (20.20.20.20)

AS_100 is now able to reach the Cloud Server, because the Attacker is successfully redirecting traffic to the right path effectively establishing the man-in-the-middle attack [11]. The MITM attack hence goes unnoticed.

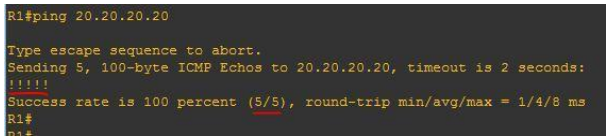

Fig 9: Shows successful pings to cloud server through AS_600

4.4 Scenario 4 – Upstream Providers Filter Traffic Using Prefix-List

All Upstream for the Attacker (AS_20 and AS_800), now begin to do proper filtering, making sure that the Attacker network, is only allowed to advertise prefix that belongs to it. The filtering is done with a prefix-list.

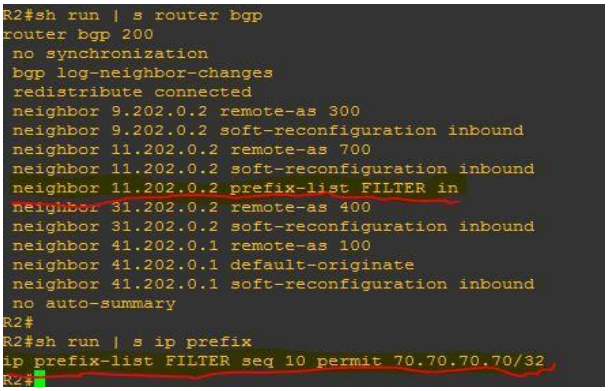

Fig 10: Shows the implementation the prefix-list

The application of the prefix list effectively stops the MITM attack by restricting. A trace from AS_100 to the Cloud Server now begins to take its original path through AS_300→AS_500 to the Cloud Server. See output below.

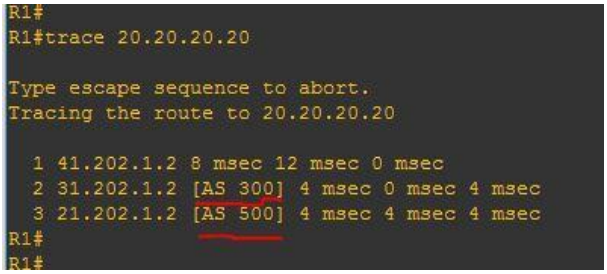

Fig 11: Shows trace output confirming the original path AS_100 takes to reach the cloud server

5. CONCLUSIONS

This study evaluates and simulates the use of prefix-lists as a viable approach to mitigating the risk of a BGP Man-In-The-Middle [11] attack in the context of cloud IaaS provider's peering relationship with its neighbor BGP routers. We believe this approach to be readily applicable in the real world sense to provide some appreciable level of confidence for cloud service providers seeking to offer high value virtualized resources largely because it is easily implementable and can be supported by the existing routing infrastructure without necessarily having to put in place hierarchical PKI to provide authentication for the identity and prefix ownership rights of ASes as may be required with sBGP, soBGP and even psBGP. This approach also successfully prevents a MITM attack instead of reporting on the attack as is the case with PHAS. Although implemented on a CISCO IOS platform, we believe that same approach could be extended and tested on JUNIPER, HUAWEI or any other enterprise grade router vendor's platform. Although capable of protecting against attack, we concede that the approach doesn't provide an automated mechanism for identifying ASes but rather may rely on the knowledge and training of network engineers in properly configuring BGP routers if the approach is to work successfully. Considering the fact that BGP attacks do not occur rampantly across the internet due to the architecture service providers employ with respect to upstream provider or peer redundancy, we conclude that the use of a prefix-list could largely suffice and by extension lead to greater availability of the cloud service.

6. REFERENCES

[1] Rashid, F. Y. (2016) Introducing the 'Treacherous 12,' the top security threats organizations face when using cloud services.http://www.infoworld.com/article/3041078/security/the-dirty-dozen-12-cloud-security-threats.html

[2] Oti, S.B. and Hayfron-Acquah, J.B., 2014. Practical Security Approaches against Border Gateway Protocol (BGP) Session Hijacking Attacks between Autonomous Systems. *Journal of Computer and Communications*, *2014*.

[3] McCullagh, D. (2008) How Pakistan knocked YouTube offline (and how to make sure it never happens again). http://www.cnet.com/news/how-pakistan-knocked-youtube-offline-and-how-to-make-sure-it-never-happens-again/

[4] Lad, M., Massey, D., Pei, D., Wu, Y., Zhang, B. and Zhang, L., 2006, August. PHAS: A Prefix Hijack Alert System. In *Usenix Security*.

[5] Kent, S., Lynn, C. and Seo, K., 2000. Secure border gateway protocol (S-BGP). *Selected Areas in Communications, IEEE Journal on*, *18*(4), pp.582-592.

[6] Butler, K.R., Farley, T.R., McDaniel, P. and Rexford, J., 2010. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, *98*(1), pp.100-122.

[7] Zhao, M., Smith, S.W. and Nicol, D.M., 2005. The performance impact of BGP security. *IEEE network*, *19*(6), pp.42-48.

[8] White, R., 2003. Securing BGP through secure origin BGP (soBGP). *Business Communications Review*, *33*(5), pp.47-53.

[9] Wan, T., Kranakis, E. and van Oorschot, P.C., 2005, February. Pretty Secure BGP, psBGP. In *NDSS*.

[10] Lad, M., Massey, D., Pei, D., Wu, Y., Zhang, B. and Zhang, L., 2006, August. PHAS: A Prefix Hijack Alert System. In *Usenix Security*.

[11] Murphy, S., 2006. BGP security vulnerabilities analysis. https://tools.ietf.org/html/rfc4272.html

[12] Empson, s., Gargano, P., Roth, H., CCNP Routing and Switching Portable Command Guide: Configuration of Redistributionhttp://www.ciscopress.com/articles/article.asp?p=2273507&seqNum=11