IJARCCE

ISSN (Online) 2278-1021
ISSN (Print) 2319 5940

International Journal of Advanced Research in Computer and Communication Engineering
ISO 3297:2007 Certified
Vol. 5, Issue 11, November 2016

Border Gateway Protocol (BGP): A Boon to Internet Application and ISPs

Vishesh S[1], Manu Srinath[1], Suhas S[2], Srivatsa S Murthy[2], Amar Tejas M[3]

B.E, Department of Telecommunication Engineering, BNM Institute of Technology, Bangalore, India[1]

Student, Department of Computer Science Engineering, BNM Institute of Technology, Bangalore, India[2]

Student, Department of Computer Science Engineering, Bangalore Institute of Technology, Bangalore, India[3]

Abstract: The internet today runs on a complex routing protocol called Border Gateway Protocol (BGP). BGP is unique compared to other routing protocols. BGP is a rich protocol that has many ways to sustain nodes or network failures as well as change in network topology, a boon to 'internetwork' or 'the internet'. In this paper, we present to you the fringe benefits that the Border Gateway Protocol offers over Internal Gateway Protocol (IGP). We discuss the major fringe benefits offered by BGP: Multihoming and Scalability, and methods of achieving the same. We also shed light on iBGP and eBGP, BGP attributes Autonomous system (AS), Autonomous System Number (ASN) and BGP route selection process.

Keywords: Border Gateway Protocol (BGP), BGP is unique, to 'internetwork', Internal Gateway Protocol (IGP), Multihoming, Scalability, iBGP, eBGP, BGP attributes, Autonomous system (AS), Autonomous System Number (ASN), BGP route selection process.

I. INTRODUCTION

Border Gateway Protocol (BGP) is a standardized Exterior Gateway Protocol (EGP), as opposed to RIP, OSPF, and EIGRP which are Interior Gateway Protocols (IGPs). BGPv4 is the current standardised deployment. BGP is considered a "Path Vector" routing protocol. BGP was not built to route within an Autonomous System (AS), but rather to route between different autonomous systems. [1] BGP maintains a separate routing table based on shortest AS path and various other attributes, as opposed to the IGP metrics like distance or cost. BGP is the routing protocol of choice on the internet. Essentially internet is a collection of interconnected Autonomous Systems. BGP Autonomous Systems are assigned an Autonomous System Number (ASN) [2] which is a 16-bit number ranging from 1-65535. A specific subset of this range, "64512 to 65535", has been reserved for private (or internal) use. BGP utilizes TCP for reliable transfer of its packets on port 179. [3] BGP is not necessary when

- Multiple connections to the internet are required
- Fault tolerance or redundancy of outbound traffic can easily be handled by an IGP, such as OSPF or EIGRP.
- If there is only one connection to an external AS (such as the internet)

BGP should be used under the following circumstances

- Multiple connections exist to external ASs (such as the internet) via different providers.
- Multiple connections exist to external ASs through the same provider, but connected via separate CO or routing policy.
- The existing routing equipment can handle the additional demands.

Figure 1 shows various ASs connected using BGP and IGP being used inside the AS. In the figure, there are four different AS- 65250, 65000, 64520 and 65500. BGP is used between different AS and IGP is used inside each AS. The routers at the border of an AS are called Border Routers and they are configured with BGP. IGP is used inside the AS (OSPF, EIGRP or any other IGPs). BGP itself has two classifications: iBGP and eBGP, called internal BGP and external BGP respectively. When BGP runs between two peers in the same AS, it is referred to as iBGP. When it runs between two ASs, it is called eBGP. The major reasons to prefer BGP over IGP are:

- Multihoming
- Scalability

And it has a rich set of attributes and policies to implement the same. [4]

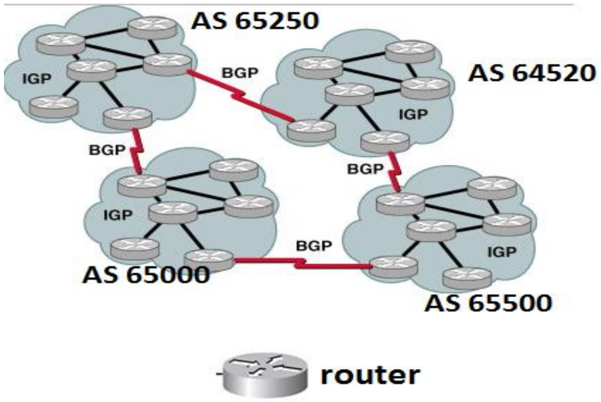

Figure 1 ASs connected using BGP and IGP being used inside the AS

DOI 10.17148/IJARCCE.2016.51109

II. BGP AND ITS FRINGE BENEFITS OVER IGP

A. Multihoming and Scalability

Multihoming is the practice of connecting a host or a computer network to more than one network. This can be done in order to increase the reliability or performance or to reduce cost. It can be achieved at many layers of the protocol stack and many places in the network.

- Multiple network interfaces in a PC.
- An ISP with multiple upstream interfaces

The question that haunts us is- Why Multihoming and how to achieve it?

1. Redundancy

One connection to the internet means the network is dependent on

- Local Router
- WAN media
- Upstream service provider

And any mishap in the above three leads to network failure. The possible mishap may be

a) Local router

- Configuration- if there is a break in the configuration, then it leads to loss of connectivity
- Software- there may be a bug in the actual OS on the router.
- Hardware- faulty hardware i.e.; faulty ICs used in the routers.

b) WAN media

Consider that the configuration is rightly done, no bug in the software and the hardware is fit, but there can be issue with WAN media and they are:

- Physical failure- this can happen when there is a break in the physical connectivity over long distance communication
- Carrier failure- carrier networks are made up of large, complex configurations of hardware, interconnected to provide communication services to people spread over large geographic areas. Any fault in this system can cause a major setback in the network.

c) Upstream ISP

- Configuration (fault in configuration)
- Software (Bug in software)
- Hardware (hardware failure)

2. Reliability

- Business critical applications demand continuous availability.
- Lack of redundancy implies lack of reliability, implies loss of revenue.

3. Supplier Diversity

Internet connection from two or more suppliers:

- With two or more diverse WAN paths
- With two or more exit points
- With two or more international connections

IGP has a limitation when it comes to the matter of scalability in the network. Whereas, BGP can deal with thousands of connections easily, which IGP cannot. BGP is currently deployed worldwide and carries approximately 155000 routing entries at the core of the internet. Some providers have been known to carry closer to 280000 routes. Policies are hard to define and enforce with an IGP because of limited flexibility. Usually a tag is the only tool available. In this age of increasingly complex networks (in both architecture and services), BGP offers an extensive suite of knobs to deal with complex policies, such as the following:

- Communities
- AS-PATH filters
- Local preference
- Multiple exit discriminator

The core is the first place in a network where scaling issues will become apparent. This is because the core tends to combine the largest number of routes with largest amount of traffic, taxing the routers to their limit. Using BGP in the core allows the routes in the core to be separated into two parts: routes within the core and routes external to the core. The iBGP mesh carries the routes external to the core, while the IGP continues to carry just the routes within the core.

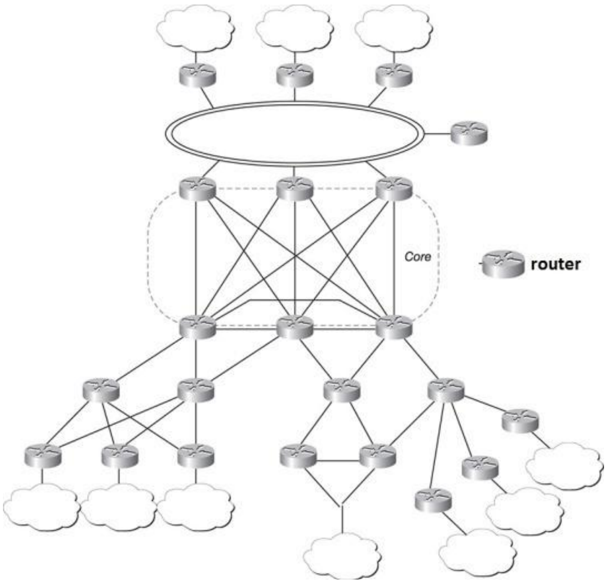

Figure 2 A typical network with a core, iBPG used in the outlying areas of network and IGP would continue to carry reachability information for connections between the routers within the core

Figure 2 shows the clear distinction between these two types of routes. iBGP would be used to carry the information about the outlying areas of the network, while the IGP would continue to carry reachability information for connections between the routers within the core itself.

DOI 10.17148/IJARCCE.2016.51109

The simplest BGP core is a simple full iBGP mesh. With this type of implementation, one can redistribute routes learned from IGP into BGP at the edge of the core, and then remove the redistributed routing information at the core edge. There are three options for removing routing information at the core edge. There are three options for removing the routing information once it has been redistributed in the BGP core. First, one could set iBGPs administrative distance lower than IGPs, so that iBGP routes are preferred over the IGP routes. This effectively filters the IGP routing information from the core. Second, one could configure a set of explicit route filters at the edge of the core. Finally, one could use a completely different IGP (or a different instance of the same IGP) in the core than the rest of the network. This approach provides an instant scalable core.

As one's network grows towards becoming an international juggernaut, taking the load of the core routers is not enough. In this case, BGP should be extended to the rest of the network. Three approaches can be followed:

- Divide the network in separate routing domains. (Connect them using eBGP)
- Use of confederations
- Use of route reflectors. [5]

Figure 3 shows dividing network into separate regions. Figure 4 shows using confederations to reduce the number of neighbours by breaking up the AS into smaller units. Now confederations make the network (of smaller pieces) look like one AS to the eBGP peers. Figure 5 shows the use of route reflectors. In short, route reflectors break the route forwarding rules of iBGP. Route reflectors can re-advertise routes learned from one iBGP peer to another iBGP peer. In figure 5, routers C and B are configured as route reflectors.

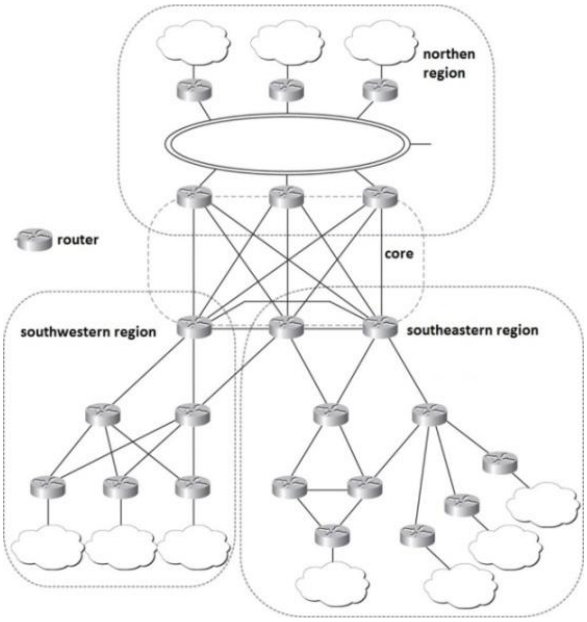

Figure 3 Divided network

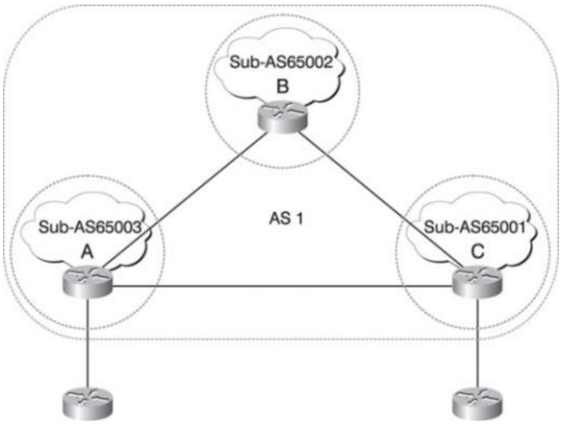

Figure 4 Use of Confederations

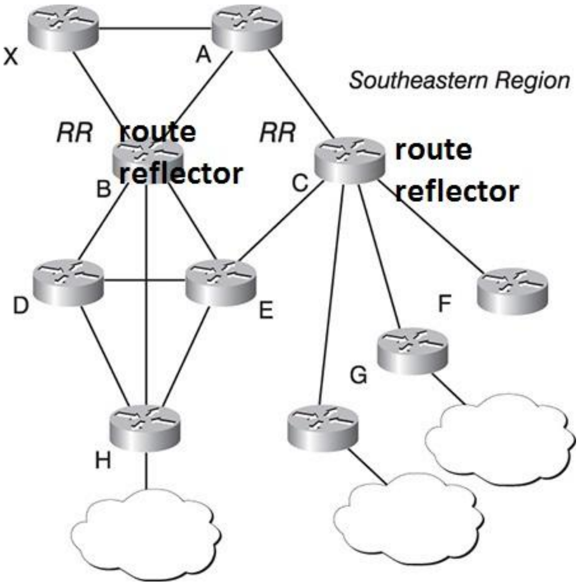

Figure 5 Use of route reflectors

B. BGP attributes

BGP utilises several attributes to determine the best path to a destination. [6] The following describes several specific BGP attributes:

- AS-path: identifies the list (or path) of traversed AS's to reach a particular destination.
- Next Hop: identifies the next hop IP address to reach a particular destination.
- Origin: identifies the originator of the route.
- Local preference: provides a preference to determine the best path for outbound traffic.
- Atomic aggregate: identifies routes that have been summarised or aggregated
- Aggregator: identifies the BGP router that performed and address aggregation
- Community: tags routes that share common characteristics into communities.
- MED: Multi-Exit Discriminator provides a preference to eBGP peers to a specific inbound router.
- Weight: similar to local preference provides a local weight to determine the best path for outbound traffic.

DOI 10.17148/IJARCCE.2016.51109

AS path, next hop, origins are well known mandatory attributes. Local preference and atomic aggregate are well known discretionary attributes. Aggregator, community are optional transitive attributes. MED is an optional non-transitive attribute. Weight is a CISCO proprietary. Each attribute is identified by a code as shown in Table 1.

Table 1 Codes for BGP attributes

Sl. No.	BGP Attribute	Code Number
1	Origin	Code 1
2	AS-path	Code 2
3	Next Hop	Code 3
4	MED	Code 4
5	Local Preference	Code 5
6	Automatic Aggregate	Code 6
7	Aggregator	Code 7
8	Community	Code 8

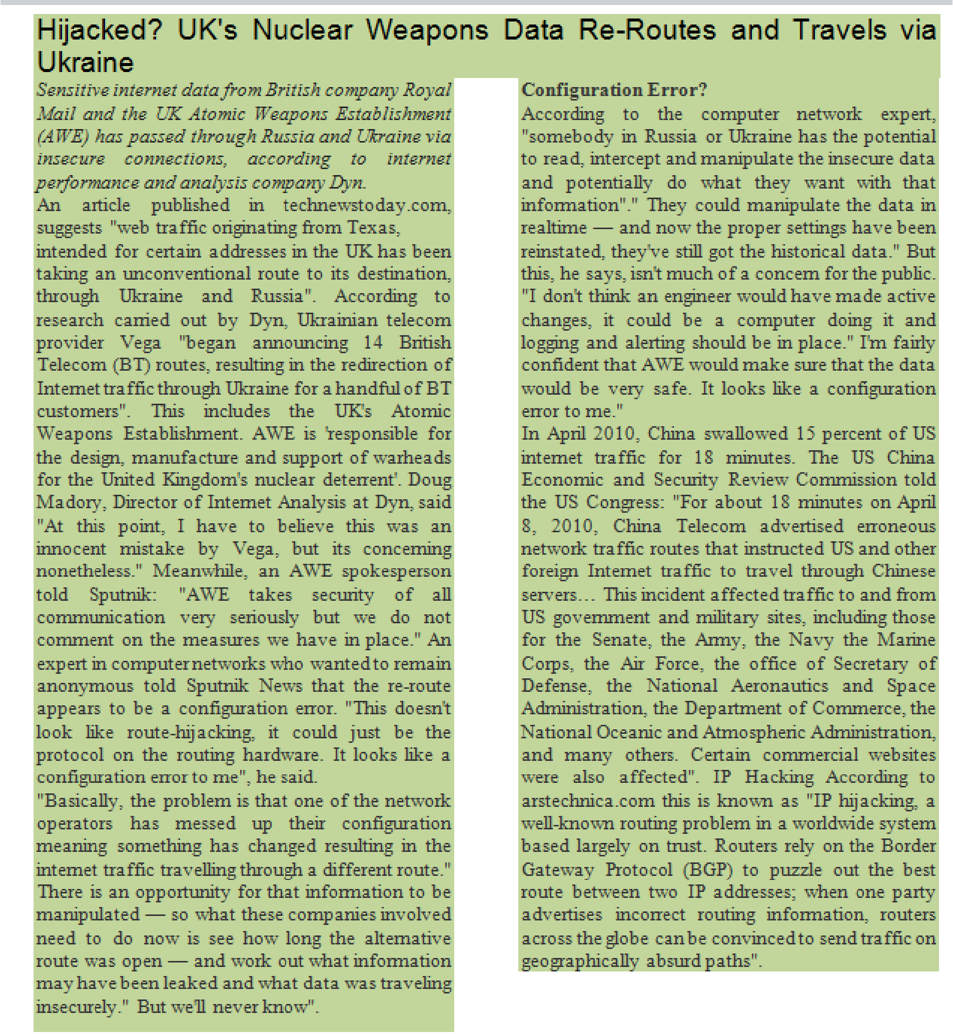

Hijacked? UK's Nuclear Weapons Data Re-Routes and Travels via Ukraine

Sensitive internet data from British company Royal Mail and the UK Atomic Weapons Establishment (AWE) has passed through Russia and Ukraine via insecure connections, according to internet performance and analysis company Dyn.

An article published in technewstoday.com, suggests "web traffic originating from Texas, intended for certain addresses in the UK has been taking an unconventional route to its destination, through Ukraine and Russia". According to research carried out by Dyn, Ukrainian telecom provider Vega "began announcing 14 British Telecom (BT) routes, resulting in the redirection of Internet traffic through Ukraine for a handful of BT customers". This includes the UK's Atomic Weapons Establishment. AWE is 'responsible for the design, manufacture and support of warheads for the United Kingdom's nuclear deterrent'. Doug Madory, Director of Internet Analysis at Dyn, said "At this point, I have to believe this was an innocent mistake by Vega, but its concerning nonetheless." Meanwhile, an AWE spokesperson told Sputnik: "AWE takes security of all communication very seriously but we do not comment on the measures we have in place." An expert in computer networks who wanted to remain anonymous told Sputnik News that the re-route appears to be a configuration error. "This doesn't look like route-hijacking, it could just be the protocol on the routing hardware. It looks like a configuration error to me", he said.

"Basically, the problem is that one of the network operators has messed up their configuration meaning something has changed resulting in the internet traffic travelling through a different route." There is an opportunity for that information to be manipulated — so what these companies involved need to do now is see how long the alternative route was open — and work out what information may have been leaked and what data was traveling insecurely." But we'll never know".

Configuration Error?

According to the computer network expert, "somebody in Russia or Ukraine has the potential to read, intercept and manipulate the insecure data and potentially do what they want with that information"." They could manipulate the data in realtime — and now the proper settings have been reinstated, they've still got the historical data." But this, he says, isn't much of a concern for the public. "I don't think an engineer would have made active changes, it could be a computer doing it and logging and alerting should be in place." I'm fairly confident that AWE would make sure that the data would be very safe. It looks like a configuration error to me."

In April 2010, China swallowed 15 percent of US internet traffic for 18 minutes. The US China Economic and Security Review Commission told the US Congress: "For about 18 minutes on April 8, 2010, China Telecom advertised erroneous network traffic routes that instructed US and other foreign Internet traffic to travel through Chinese servers… This incident affected traffic to and from US government and military sites, including those for the Senate, the Army, the Navy the Marine Corps, the Air Force, the office of Secretary of Defense, the National Aeronautics and Space Administration, the Department of Commerce, the National Oceanic and Atmospheric Administration, and many others. Certain commercial websites were also affected". IP Hacking According to arstechnica.com this is known as "IP hijacking, a well-known routing problem in a worldwide system based largely on trust. Routers rely on the Border Gateway Protocol (BGP) to puzzle out the best route between two IP addresses; when one party advertises incorrect routing information, routers across the globe can be convinced to send traffic on geographically absurd paths".

Figure 6 Research on internet performance and analysis by the company Dyn on “Sensitive internet data from British Royal Mail and the UK Atomic Weapons Establishment (AWE) has passed through Russia and Ukraine via insecure connection.”

DOI 10.17148/IJARCCE.2016.51109

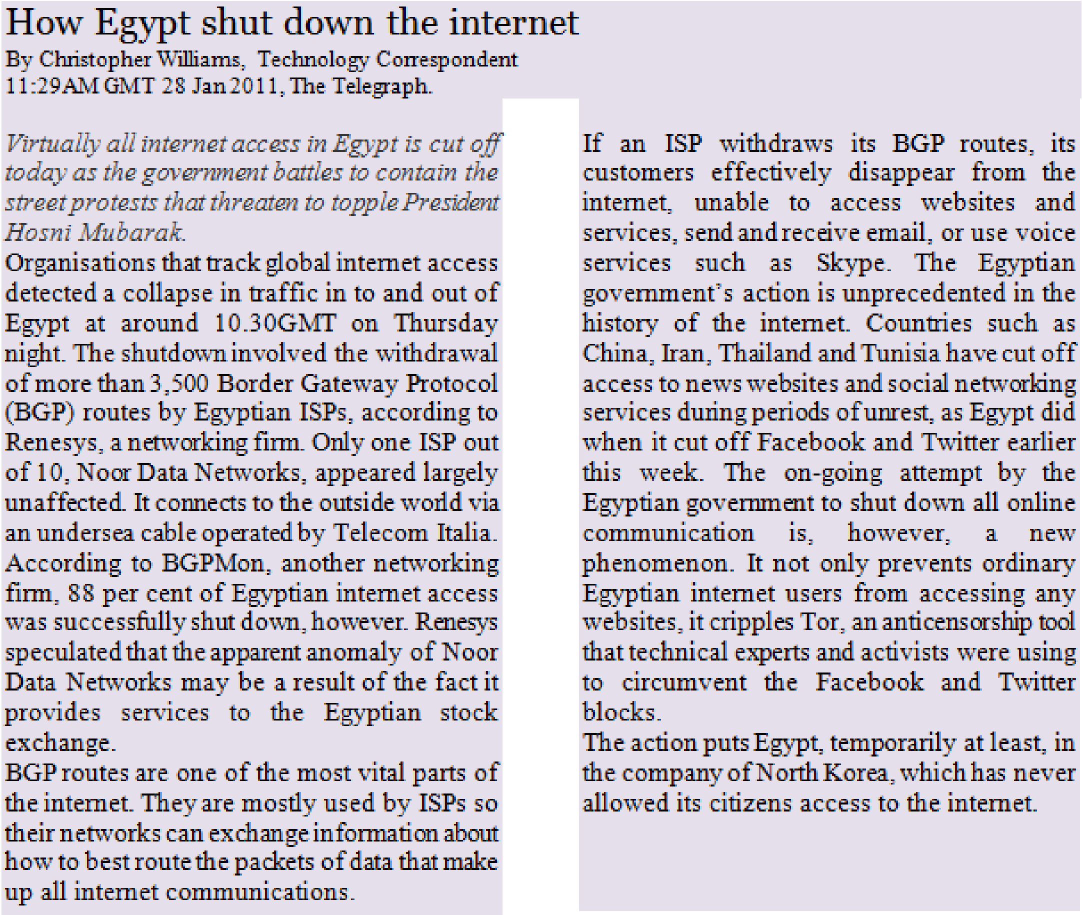

How Egypt shut down the internet

By Christopher Williams, Technology Correspondent
11:29AM GMT 28 Jan 2011, The Telegraph.

Virtually all internet access in Egypt is cut off today as the government battles to contain the street protests that threaten to topple President Hosni Mubarak.

Organisations that track global internet access detected a collapse in traffic in to and out of Egypt at around 10.30GMT on Thursday night. The shutdown involved the withdrawal of more than 3,500 Border Gateway Protocol (BGP) routes by Egyptian ISPs, according to Renesys, a networking firm. Only one ISP out of 10, Noor Data Networks, appeared largely unaffected. It connects to the outside world via an undersea cable operated by Telecom Italia.

According to BGPMon, another networking firm, 88 per cent of Egyptian internet access was successfully shut down, however. Renesys speculated that the apparent anomaly of Noor Data Networks may be a result of the fact it provides services to the Egyptian stock exchange.

BGP routes are one of the most vital parts of the internet. They are mostly used by ISPs so their networks can exchange information about how to best route the packets of data that make up all internet communications.

If an ISP withdraws its BGP routes, its customers effectively disappear from the internet, unable to access websites and services, send and receive email, or use voice services such as Skype. The Egyptian government's action is unprecedented in the history of the internet. Countries such as China, Iran, Thailand and Tunisia have cut off access to news websites and social networking services during periods of unrest, as Egypt did when it cut off Facebook and Twitter earlier this week. The on-going attempt by the Egyptian government to shut down all online communication is, however, a new phenomenon. It not only prevents ordinary Egyptian internet users from accessing any websites, it cripples Tor, an anticensorship tool that technical experts and activists were using to circumvent the Facebook and Twitter blocks.

The action puts Egypt, temporarily at least, in the company of North Korea, which has never allowed its citizens access to the internet.

Figure 7 "How Egypt shut down the internet" The Telegraph, 28th January, 2011

C. BGP "Best Path" Determination

If BGP contains multiple routes to the same destination, it compares the routes in pairs, starting with the newest entries and working towards the oldest entries.

BGP determines the best path by successively comparing the attributes of each "Route pair". The attributes are compared in a specific order: [7]

- Weight: which route has the highest weight?
- Local preference: which route has the highest route preference?
- Locally originated: did the local router originate this route? In other words, is the next hop to the destination 0.0.0.0?
- AS-path: which route has the shortest AS-path?
- Origin code: where did the route originate?
- MED: which path has the lowest MED?
- BGP route type: is it eBGP or iBGP route? eBGP routes are preferred.
- Age: which route is the oldest? Oldest is preferred.
- Router ID: which route originated from the router with the lowest BGP router ID?
- Peer IP address: which route originated from the router with the lowest IP?

III.DISCUSSIONS

BGP can be attacked in many ways. Communication between BGP peers can be subjected to active or passive wiretapping. The BGP software, configuration information, or routing databases of a router may be modified or replaced via unauthorized access to a router, or to a server or management workstation from which router software is downloaded. These latter attacks transform routers into hostile insiders, so security measures must address such Byzantine failures. Figure 6 shows research on internet performance and analysis by the company Dyn on "Sensitive internet data from British Royal Mail and the UK Atomic Weapons Establishment (AWE) has passed through Russia and Ukraine via insecure connection." [8] Improved physical and procedural security for network management facilities, and routers, and cryptographic security for BGP traffic between routers would help reduce some of these vulnerabilities. However, physical and procedural security is expensive and imperfect, and these countermeasures would not protect the Internet against accidental or malicious misconfiguration by operators, nor against attacks that mimic such errors. Misconfiguration of this sort has been a source of Internet outages in the past and seems likely to persist. Any security approach that relies on ISPs to act properly violates the "principle of least privilege" and leaves the Internet routing system vulnerable at its weakest link. Routers also are susceptible to resource exhaustion attacks based on delivery of large quantities of management traffic, BGP or otherwise. This vulnerability arises because these devices are designed with the not unreasonable model that management traffic is a very tiny percentage of all the traffic that arrives at a router. Router interfaces can deliver traffic to the management processor at very high rates, because they are designed to accommodate subscriber traffic flows. As usual, there are merits and demerits in every protocol/system, which is eliminated by new updates in the present technology.

IV. CONCLUSION

There are several protocols that are used for connectivity and routing inside an Autonomous System such as OSPF, EIGRP, IS-IS, etc. The question that arises is: What if several autonomous systems are to be connected? Can IGP protocols like OSPF, EIGRP, IS-IS have the capability to deal with thousands of connections, i.e. are they scalable to such high demands? The solution to this is the Border Gateway Protocol. It has a very high scalability and Multihoming can be achieved by using rich set of policies and attributes that the BGP offers. Multihoming can bring reliability and can enhance performance in a network. Figure 7 explains how without Multihoming, i.e. by not creating redundancy, the real time problems faced when an ISP withdraws its BGP routes, its customers effectively disappear from the internet, unable to access websites and services, send and receive email or use voice services such as Skype. [9] This move by the Egyptian government not only prevented ordinary citizens from accessing any websites, it crippled 'TOR', an anti-censorship tool that technical experts and activists were using to circumvent the Facebook and Twitter blocks.

REFERENCES

[1] Orbit-Computer-Solutions.Com(n.d), Computer Training & CCNA Networking Solutions, Orbit-Computer-Solutions.com, retrieved 8 October 2013, <"Archived copy". Archived from the original on 2013-09-28. Retrieved 2013-10-08.>

[2] Hawkinson & Bates, Best Current Practice, Guidelines for creation, selection, and registration of an Autonomous System (AS).

[3] 'Port 179 Details' http://www.speedguide.net/port.php?port=179

[4] Free Cisco labs for CCNA, CCNP and CCIE students! Presented by René Molenaar - CCIE #41726

[5] 'BGP Route Reflector' https://networklessons.com/bgp/bgp-route-reflector/

[6] 'BGP attributes' https://networklessons.com/cisco/ccie-routing-switching/

[7] 'bgp best path determination' https://nets.ucar.edu/nets/devices/routers/bgp/bgp-route-selection-algorithm.shtml

[8] Research on internet performance and analysis by the company Dyn on "Sensitive internet data from British Royal Mail and the UK Atomic Weapons Establishment (AWE) has passed through Russia and Ukraine via insecure connection."- https:// http://sputniknews.com/science/201503171019620104/

[9] "How Egypt shut down the internet" The Telegraph, 28th January, 2011

BIOGRAPHIES

Vishesh S who hails from Bangalore (Karnataka) has completed B.E in Telecommunication Engineering from VTU, Belgaum, Karnataka in 2015. His research interest includes embedded systems, wireless communication and medical electronics. He is also the founder of the firm 'Konigtronics'.

Manu Srinath who hails from Bangalore (Karnataka) has completed B.E in Telecommunication Engineering from VTU, Belgaum, Karnataka. His research interests include networking, image processing and cryptography. He is currently working as a Design Engineer in the start-up "Konigtronics".

Suhas S hails from Bangalore (Karnataka). He is currently pursuing B.E in Computer Science Engineering at BNM Institute of Technology. His research interests include Cloud Computing and Software Engineering and Management.

Srivatsa S Murthy hails from Bangalore (Karnataka). He is currently pursuing B.E in Computer Science Engineering at BNM Institute of Technology.

Amar Tejas M hails from Bangalore (Karnataka). He is currently pursuing B.E in Computer Science Engineering at Bangalore Institute of Technology. His research interests include Database Management Systems and Software Engineering.

DOI 10.17148/IJARCCE.2016.51109