

BORDER GATEWAY PROTOCOL

BORDER GATEWAY PROTOCOL FILTERING GUIDELINES

MARCH 2008

Freedom of Information Act (FOIA)

This information is supplied in confidence and may not be disclosed other than to the agreed readership, without prior reference to CPNI. Within the UK, this material is exempt from disclosure under the relevant Freedom of Information Acts and may be subject to exemption under the Environmental Information Regulations and the Data Protection Act 1998.

Disclaimer

Reference to any specific commercial product, process or service by trade name, trademark manufacturer, or otherwise, does not constitute or imply its endorsement, recommendation, or favouring by CPNI. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes

CPNI has taken every care in preparing this protective security advice, which is informed by intelligence on the threat. However, CPNI cannot accept any liability to any person or company for any financial loss or damage arising from the use of this advice or from any failure to give advice.

CPNI acknowledges the assistance of members of the Network Security Information Exchange (NSIE) in providing the input into this process, and wishes to thank them for their efforts and co-operation.

Document History	Version	Date of Issue
Updated CPNI version	02	March 2008
Original NISCC version	01	April 2004

BORDER GATEWAY PROTOCOL FILTERING GUIDELINES

Summary of Principal Risks

1. Router failure or impaired performance.
2. Blackholing internet traffic to some addresses.
3. Isolation of internet networks or subnets.
4. Traffic misdirection/route hijack.
5. Eavesdropping.

All of these have cost implications for the providers, as well as CNI security implications, and all also carry a risk of loss of reputation/customer confidence.

Threats

May be generally categorised as accidents, insiders, attack via a weaker peer or collateral damage from other activity.

1. Many of these effects can be caused by accidental mis-configuration, either in the core network, or a directly or indirectly connected network.
2. There is a threat from a subverted disaffected or malicious system administrator, who could deliberately carry out routing mis-configuration, in an effort to attack either this network, or a network in another Autonomous System (AS).
3. A privileged individual in another AS (customer or peer) could use their position to launch an attack (denial of service) against the AS or against the internet generally.
4. A network-based attack from elsewhere on the internet, for example side effects of virus attacks causing excess BGP traffic, due to increased traffic loads. A similar effect can be seen from some other IP attacks.
5. Deliberate BGP attack via a poorly configured ISP from elsewhere on the Internet.

6. Network-based attack on one or more routers within the AS (whether focussed, e.g. telnet/SNMP, or unfocussed e.g. SYN flood) is not directly a BGP vulnerability, but by creating instability within the AS may give rise to route flapping as seen from outside. This in turn may cause BGP instability. Also, e.g. flooding attack on BGP TCP port could disable BGP, though this is really a TCP/IP attack.

Guidelines

These guidelines identify generally accepted practices for Border Gateway Protocol (BGP) filtering. Readers may also find this document useful in determining guidelines to consider when implementing BGP. These practices are one of many parts of running a robust and healthy system. As such, the implementation of these practices must be viewed in the context of the whole system. Other router threats can bring significant instability to routing. Abatement of these risks is beyond the scope of this document; however, it must be considered in addition to these guidelines. Implementation of the practices identified must be done in a manner that would not cause additional risk of instability, excess load, or increased risk.

Generally Accepted Practices

- ▶ Deny special prefixes assigned and reserved for future use
- ▶ Deny unallocated (grey/bogon) space
- ▶ Deny over-specific prefix lengths
- ▶ Maintain route flap dampening default settings or set according to RIPE parameters
- ▶ Aggregate routes where possible
- ▶ Deny exchange point prefixes
- ▶ Deny routes to internal IP spaces
- ▶ Minimize BGP use with customers¹ (e.g. single homed customers, customers who do not themselves have downstream customers)²
- ▶ Restrict routes exchanged with customers to those concerning “customer-declared” IP space
- ▶ Restrict routes exchanged with peers, depending on relationship with peers and between peers³

Implementation of these practices should be considered in the context of existing relationships with customers and business partners and customer requirements.

¹ For purposes of this paper, customers are not generally end users.

² Excluding use of BGP for customer provided remote triggered blackhole filtering.

³ This may not be feasible depending on the size of the peers table and access to correct Autonomous System/route-sets among peers.

Other Filtering Practices

The following guidelines are presented for consideration as additional practices to improve BGP security. Implementation of such guidelines would depend heavily on existing relationships with customers and routing partners as well as other routing arrangements. In addition, there may be technical or business issues associated with their implementation. As with all security decisions, the benefits and costs of implementation should be weighed against one another.

- ▶ Deny over-general prefix lengths
- ▶ Deny inappropriate length RIR allocations
- ▶ Deny inappropriate announcements for legacy A/B/C space
- ▶ Protect critical networks by defining three levels of networks (e.g. platinum, gold, silver) and only allowing certain routes on each level⁴
- ▶ Deny route flap dampening on “golden networks”⁵
- ▶ Implement BGP graceful restart (e.g. Cisco non-stop forwarding or other Cisco GRIP features)⁶

Authentication of BGP Sessions

Use MD5 authentication on peering links wherever practical, using appropriate password strength (which may depend on customer requirements).

Vulnerability Management

As with any other IT-based system, BGP requires a robust process for handling vendor’s patches/workarounds for reported vulnerabilities.

⁴ General consensus is that the costs/risks associated with this practice may outweigh the benefit.

⁵ General consensus is that network operators should address route dampening.

⁶ General consensus is that the costs and complexity introduced by this guideline may greatly exceed the risk associated with BGP vulnerabilities.

ANNEX A:

Additional information which may be useful for Providers

Checklist of baseline security practice (UK implementation)

Note: This aligns to the generally accepted practices above, the only exception being “deny short prefixes (</6)”, which may be difficult for some US providers, if they are advertising several adjacent Class A addresses, but was agreed by all UK providers to be valuable in their environments.

		Implemented?
	Filtering:	
1	Deny assigned special prefixes	
2	Deny special prefixes reserved for future use	
3	Deny unallocated prefixes (grey space/bogons)	
4	Deny prefixes > /28	
5	Deny prefixes < /6	
6	Deny exchange point peering mesh prefixes	
7	Deny routes to internal IP space	
	Configuration & Management:	
8	Implement route flap dampening using RIPE recommended parameters (generally default parameters in router)	
9	Minimize BGP use with customers (e.g. single homed customers, customers who do not themselves have downstream customers)	
10	Restrict routes accepted from customers to those concerning customers declared IP space	
11	Only exchange own and customer routes with peers (depending on relationship with & between peers)	
12	Use shortest possible prefix to exchange own & customer routes	
13	Aggregate routes where possible, (suggested: all routes more specific than /20)	
	Authentication:	
14	Use MD5 authentication on peering links wherever practical, using appropriate password strength (which may depend on customer requirements)	

	Vulnerability Management:	
15	Have a robust process for applying manufacturers patches/workarounds for reported vulnerabilities	

Checklist of additional suggested security measures

		Implemented?
16	Deny inappropriate length RIR allocations	
17	Deny inappropriate announcements for legacy A/B/C space	
18	Monitor announcements with "borderline" length prefixes (e.g. /6 /28)	
19	Set max-prefix limits on IXP and customer peerings	
20	Implement BGP graceful restart, for example Cisco Nonstop forwarding or other Cisco GRIP features	

ANNEX B:

Other suggestions for improving BGP security – under discussion in the community.

1. [12] MD5 password strength & management issues. There are a variety of recommendations for best practice in MD5 use. It is certainly true that MD5 cracking tools are available in the hacker community. Currently, most providers take the pragmatic view that MD5 implementation (with well-chosen, non dictionary-based passwords) enhances security, even where some details of the implementation could be further improved.
2. Egress filtering. Ideally, filtering should be both on ingress & egress. Providers should encourage their customers to perform egress filtering, according to these guidelines. Outbound route filtering can provide an alternative to customer egress filtering.
3. [14] The Generalized TTL Security Mechanism (GTSM) – previously known as the BGP TTLH (Time To Live Hack) or the BTSH (BGP Time-to-live Security Hack) – suggests that, where the peer is always exactly one hop away, routers set TTL to be 255, and only accept packets with TTL 254. This introduces extra difficulty for an attacker, compared to the default of expecting the TTL to be 1. For BGP, GTSM could be used to protect against DOS attacks flooding port 179. This is not widely implemented (Juniper JUNOS implements this feature).
4. Filtering of incoming BGP packets on entry, before passing to the CPU, can also protect the router from DOS via port 179. Implementation is not even across vendors, feedback from providers on which implementations are most successful would be welcome.
5. [9] [11] Graceful restart. A facility where a router preserves its forwarding state through a restart (time-limited to e.g. 30 seconds), so eliminating the need for peers to “flap” all its routes. Implemented by Juniper, Cisco (as Cisco non-stop forwarding) & Riverstone, among others. Several providers (US) suggest that the cost of implementing this feature outweighs the benefit.
6. [18] Sink-holes to protect customer networks. The ability to set these up in response to an event/attack is highly desirable, and will provide a valuable service to CNI providers in non-communications sectors.
7. AS-Path filtering. In addition to prefix filtering, some providers recommend using AS-Path filtering, to drop any announcements with private AS numbers in the AS path, and setting a max-as-path-length.
8. [13] Use of BGP MED values has been suggested as a means to enhance security; however, this can have a variety of effects. Accepting MED values from peers may introduce a (small) extra risk.
9. [15] S-BGP aims to set up a PKI authentication scheme to authenticate BGP peering & announcements. Lack of backward compatibility with BGPv4 is a problem.
10. [16, 17] soBGP. Developed by Cisco. Aims to use PKI to authenticate the origin of BGP packets, but not the peering connections. Security increases proportionally with adoption, but it is compatible with non-secure BGPv4.

REFERENCES

- [1] "Routing TCP/IP, Volume II", Doyle & Carroll, Cisco Press
- [2] RFC 3330: Special-Use IPv4 Addresses
<http://www.ietf.org/rfc/rfc3330.txt>
- [3] Cisco suggested filters: T-ip-prefix-filter-ingress-strict-check.txt
<http://www.cisco.com/public/cons/isp/security/T-ip-prefix-filter-ingress-loose-check.txt>
- [4] The Team Cymru Bogon List
<http://www.cymru.com/Documents/bogon-list.html>
- [5] List of Golden Networks
<http://www.cymru.com/gillsr/documents/golden-networks>
- [6] RIPE Routing WG Recommendations for Co-ordinated Route-flap Dampening Parameters, Panigl, Schmitz, Smith, Vistoll
- [7] "Understanding BGP Misconfiguration", Mahajan, Wetherall & Anderson
<http://www.sigcomm.org/sigcomm2002/papers/bgpmisconfig.pdf>
- [8] "Secure BGP template V. 2.4", Rob Thomas
<http://www.cymru.com/Documents/secure-bgp-template.html>
- [9] RFC 4724: Graceful Restart Mechanism for BGP, S. Sangli et al
<http://www.ietf.org/rfc/rfc4724.txt>
- [10] NRIC best practice guidelines
<http://www.bell-labs.com/cqi-user/krauscher/bestp.pl> & <http://www.nric.org>
- [11] Riverstone networks BGP support pages
http://www.riverstonenet.com/technology/bgp_restart.shtml & others
- [12] RFC 3562: Key Management Considerations for the TCP MD5 Signature Option, M. Leech
<http://www.ietf.org/rfc/rfc3562.txt>
- [13] RFC 4451: BGP MULTI_EXIT_DISC (MED) Considerations, D. McPherson
<http://www.ietf.org/rfc/rfc4451.txt>
- [14] RFC3682: The Generalised TTL Security Mechanism (GTSM)
<http://www.ietf.org/rfc/rfc3682.txt>
- [15] S-BGP internet draft, C. Lynn et al.
<http://www.ietf.org/internet-drafts/draft-clynn-s-bgp-protocol-01.txt>
- [16] Deployment Considerations for Secure Origin BGP (soBGP), Russ White, Cisco
<http://www.ietf.org/internet-drafts/draft-white-sobgp-bgp-deployment-01.txt>
- [17] Extensions to BGP to Support Secure Origin BGP (soBGP), James Ng, Cisco
<http://www.ietf.org/internet-drafts/draft-ng-sobgp-bgp-extensions-01.txt>
- [18] Deploying & Using Sinkholes – NANOG presentation by Barry Greene, Cisco & Danny McPherson, Arbor
<http://www.nanog.org/mtg-0306/sink.html>