National Cyber Security Centre
a part of GCHQ

TECHNICAL REPORT:

RESPONSIBLE USE OF THE BORDER GATEWAY PROTOCOL (BGP) FOR ISP INTERWORKING

Version 1.0

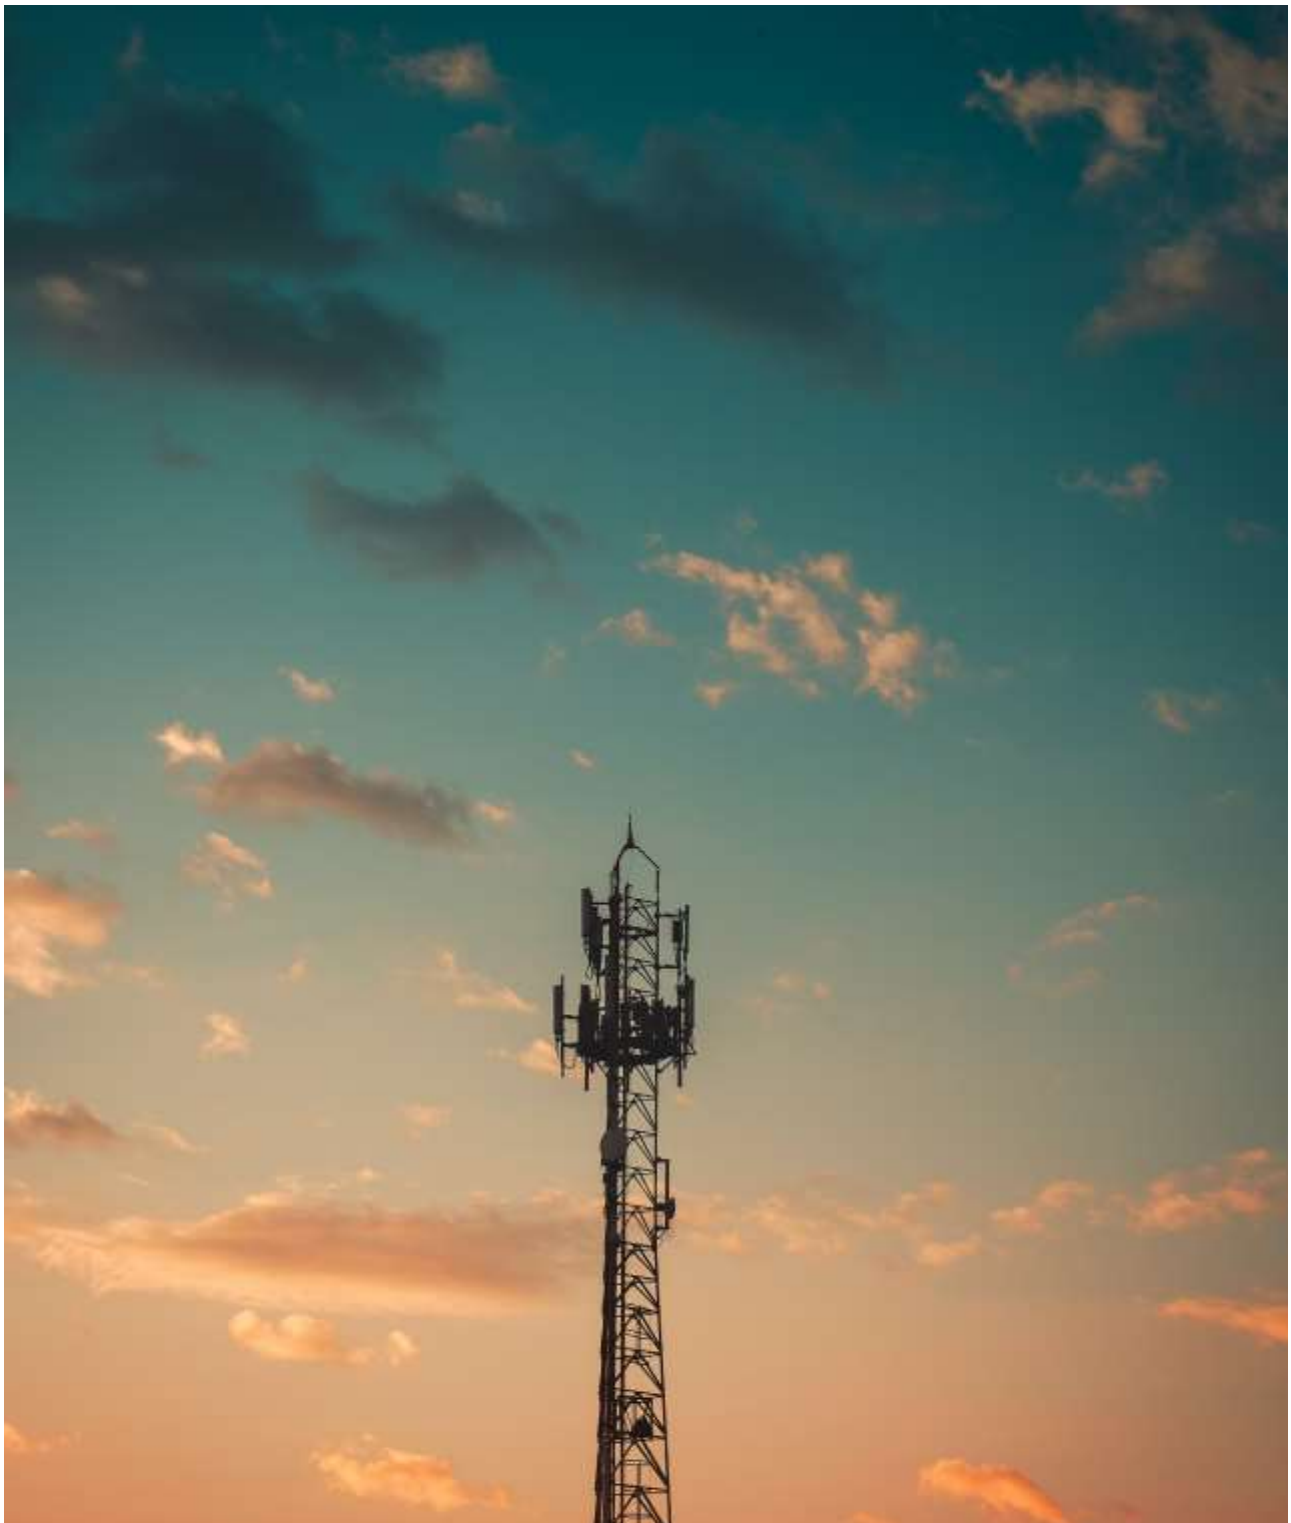

Contents

INTRODUCTION

1. This document outlines and consolidates best practice in terms of BGP (Border Gateway Protocol) Security. Whilst initially written with the UK ISP (Internet Service Provider) community in mind, the contents of this document and the principles on which it is built are broadly applicable in all BGP deployments, globally.

2. This document assumes that the reader has a good working knowledge of networking, TCP/IP and associated protocols and mechanisms including BGP and associated routing control methodologies.

3. Where this document refers to RIPE (https://www.ripe.net/), it is assumed that this also includes any other RIR (Regional Internet Registries).

4. Where the terms “shall” and “should” are used in this document, they are as defined within RFC2119.

5. The *CPNI BGP good practice guide* [reference 1] was published in 2011, with the objective of providing a single document which could be used to enhance the quality of BGP deployment, configuration and operations.

6. The original CPNI document recognised the correct implementation and operation of BGP, together with the health and housekeeping of peering arrangements as a critical element of running successful and reliable networks, both public and private.

7. Since the original document was published, the importance of the stability of IP based networks and the routing information that underpins the delivery of protocols and services has increased dramatically. This has increased the requirement for safe and secure methodologies and practices to secure BGP and maintain the integrity of the routing data that it supports and controls the distribution of.

8. As the importance of “trusted” BGP services has increased, several industry wide initiatives and recommendations have been developed. These “best practices” are continually refined and updated by the various members of the communities as issues are discovered and resolved.

9. To ensure that the UK ISP community adopt the latest best practice appropriate to their BGP deployments, reference must be made, as a minimum, to the current versions of references 2, 3, 4 and 5 at the end of this document. It should be noted that this document makes reference to these solely for further background and for explanation of the recommendations stated within this document.

10. This document will identify the key items that UK ISPs should address to ensure that they are recognised as being leaders in the adoption of BGP best practice with the implementation requirements / methodologies from references 2, 3, 4 and 5.

Guiding Principles

11. The key objectives of the NCSC and UK ISP community are to ensure that Routing Information originating from the community is as accurate and secure as possible. As far as practical, this limits the scope for accidental or malicious misuse.

12. Minimum implementation

To help ensure that Routing Information originating from the community is as accurate and secure as possible, Service Providers shall, at a minimum, implement several basic elements. These are:

12.1 Service Providers shall ensure that contact details are current and accurate on all the recognised registries, e.g. Regional Internet Registries (RIPE, APNIC etc.) and other useful locations, such as Peering DB. Note that all appropriate fields and record types should be secured appropriately, to prevent misuse.

12.2 All address space allocated to a Service Provider shall be correctly recorded in such a way that it is simple to identify and contact the "owner" to assist in resolving issues.

12.3 BGP Filtering. Service Providers shall implement ingress and egress route filtering.

To make this more effective, Service Providers shall adopt and implement mechanisms that prevent IP address spoofing, including the "BCP38" recommendations. It is advised that adherence to this is checked at https://www.caida.org/projects/spoofer/[1].

12.4 BGP monitoring. Service Providers should have a monitoring capability and actively use it to detect and monitor incidents, including (but not limited to) hijacking and denial of service attacks.

13 Service Providers should utilise BGP prefix origin validation using Resource Public Key Infrastructure (RPKI).

BGP prefix origin validation reduces the ability of third parties to accidently or deliberately 'hijack' BGP routes. The owner of a BGP prefix specifies the correct source AS (Autonomous Systems) and the prefix length(s) that are permitted to be routed on the Internet, allowing an organisation implementing origin validation filtering to use this information to reject any prefix of incorrect length or origin AS, reducing the effectiveness of route hijacking.

There are two parts to implementing origin route validation:

13.1 Route Signing:

Any LIR (Local Internet Registry) interested in protecting their network should sign their prefixes to accurately match their routing requirements. This includes ISP's infrastructure and customer prefixes, but also corporate networks with their own LIR. Once routes are signed, the LIR must not announce prefixes that do not match their signed routes. For example, temporarily announcing a more specific route to manage traffic, or temporarily

[1] Note that data utilised at this site is publicly generated and may not be fully representative of the Service Provider's network. Service Providers should deploy their own instance of the tool.

announcing from a different source AS to manage a DDOS attack, as this route would be rejected by any ISP filtering against route origin information.

PI (Provider Independent) End User resources allocated by RIPE should also be signed, either by the end user themselves, their agent, or the Sponsoring LIR, depending on who maintains the organisation object within the RIPE Database.

Legacy address resources are considered out of scope for this document owing to the lack of a centralised signing authority.

13.2 Route Filtering:

Filtering of routes is most important for networks with multi-party peering arrangements. A network served by two ISPs that already perform origin verification would see no benefit.

As route filtering based on prefix origin carries risks, it can be implemented in two phases to gain confidence.

Phase1 - monitor mode. Report prefix information against three categories of routes seen in the routing table:

- RPKI registered, has valid AS and is not longer than allowed
- RPKI registered and seen from wrong AS or longer than allowed
- Not RPKI registered

Over time, analysis of the reports generated should give confidence that routes which are both registered and invalid (wrong AS and/or longer prefix length) are genuinely bad routes that can be safely dropped.

Phase2 - full filtering. The required actions for each route are:

- RPKI registered, has valid AS and is not longer than allowed – *Accept, subject to other BGP route filters.*
- RPKI registered and seen from wrong AS or longer than allowed – *Reject*
- Not RPKI registered – *Accept, subject to other BGP route filters*.

RECOMMENDED USE OF BGP

14 The table below outlines the key areas to be considered and addressed via "best practice" when designing, deploying, operating and maintaining BGP and the associated Peering Relationships.

BGP Security		**Standards and Best Practices**	
Feature Type	Target Area	Recommendations	Key Reference Documents
Filtering	General Prefix	Filter: • Special Use Addressing • BOGONs (although RFC 6441 should be considered) • Over-specific prefix lengths • Own prefixes • Own AS • IXP LANs	NIST SP800-54 Section 4.2.1-3 RFC 1918 RFC 5735 RFC 6598 RFC 6441 BCP 194 6.1.5.1 It is suggested that the following are informational resources that may be consulted: Team Cymru Bogon list
	Advanced Prefix	For IPv4, all prefixes are allocated, so there is no reason to check for IANA-allocation. However, check that the IANA-allocation of IPv6 address space is valid. ROA (Route Origin Authorisation) of updates is recommended, e.g. RPKI. Refer to 1 above for more information on this topic. Transit providers shall implement checks on RIR database before advertising (and accepting) routes. Maximum prefix tracking (at least to a warning level) is recommended in order to identify advertisement scale-based attacks.	RFC7454 6.1.2.1 RFC7454 6.1.2.2 (RIR-Allocated Prefix Filters) RFC 6480 (SIDR) + RFC 6811 (Origin validation)

BGP Security		Standards and Best Practices	
Feature Type	Target Area	Recommendations	Key Reference Documents
	AS_PATH	Relative lengths of AS_PATH should be checked for validity. Some degree of checking is required to ensure an excessively long AS-PATH is not accepted. It is challenging to define a definite value as it will vary according to network. As standard, you should not receive your own AS or own prefixes. Note that there may be exceptions for large ISPs, e.g. for DDoS mitigation.	
	Egress (outbound to the internet)	Appropriate outbound filtering should be applied in alignment with inbound filtering rules (i.e. when exporting, peers should honour what they deem as invalid when importing).	
	Forwarding (ingress from customer, egress to the internet)	Appropriate inbound and outbound filtering should be applied in alignment with the guiding principles (See section II) to ensure invalid AS_PATH / prefixes are not forwarded. In addition to routing filters, data path filters (i.e. ACLs) should be used to enforce valid IP data sources.	
	Alignment	Ingress filters of a peer should be aligned with egress filters of its opposite peer.	NIST SP800-54 Section 4.2
	Off-router monitoring	A BGP monitoring platform should be deployed to monitor for competing announcements for your own address range(s), or for significant changes in the announcements from specific peers (e.g. AS_PATH changes, prefix count changes etc.).	CSRIC 5.2.1.5 BGP Spotlight tool (refer to NCSC)
Scalability	Prefix-limits	Prefix-limits should be applied where expected prefix count can be known and managed for peers.	CSRIC 5.4.1.2.1

BGP Security		Standards and Best Practices	
Feature Type	Target Area	Recommendations	Key Reference Documents
	Aggregation	Aggregate routes where possible and advertise covering summary prefixes based on blackhole routes (i.e. route to null) to avoid updates for inactive aggregates. Advertisement of prefixes more specific than those normally accepted should have consideration made to limit onward propagation. For example, by use of NO-EXPORT, NOPEER or provider-specific communities.	RIPE-399 RIPE-532 RFC-1197 RFC-3765
	Control Plane Policing	Mechanisms to prevent CPU/RP processing of irrelevant packets (e.g. non-configured protocols on the interface) should be applied to limit the possibility of resource exhaustion attacks. Drop policies for 'out-of-profile' (but valid) protocol packets should be deployed as part of Control Plane Policing per peer, to limit the ability of a single peer sending massive protocol updates, preventing processing of valid updates from other peers.	CSRIC 5.1.2.2.4
Stability	Route-flap Dampening	Current recommendation is that it **should not** be used. If used, careful consideration to damping settings per prefix-length must be made, along with careful consideration of penalties (i.e. half-life etc.) and whether any whitelisted (or "Golden") prefixes be excluded from RFD (this may include DNS root/gtld servers and/or other important prefixes).	RIPE 178/210/229 were obsoleted by RIPE 378, which was obsoleted by RIPE 580 (current). NIST SP800-54 recommends against the use unless strong reasons exist in favour of it.

BGP Security		Standards and Best Practices	
Feature Type	Target Area	Recommendations	Key Reference Documents
	Graceful Restart	Whilst of less value in a largely NSR (Non-Stop Routing) enabled control plane environment, graceful restart remains beneficial in some scenarios and may optionally be applied. **Note**: Extreme care must be taken when using BFD in conjunction with BGP Graceful Restart.	
Security	Generalised TTL Security Mechanism (GTSM)	Where available, RFC 3682 should be configured between peers to limit the radius of valid BGP senders. GTSM also limits the possibilities for a hash calculation 'burden' attack being used.	CSRIC 5.1.1.1 CSRIC 5.1.2.2.2 NIST SP800-54 4.4
	Authentication (MD5)	MD5 authentication (despite being deprecated as an algorithm) should be used for protocol sessions, noting that this mechanism is useful not only for authentication but also for validation of configuration (i.e. misconfiguration of peering can be expected to fail MD5). BCP194 calls out TCP-AO as recommended where available.	NIST SP800-54 4.5 RFC 5925
	MD5 Key Management	An appropriate MD5 key management and exchange mechanism must be in place between peers, to include retention and renewal policy.	

BGP Security		Standards and Best Practices	
Feature Type	Target Area	Recommendations	Key Reference Documents
	ACLs	Permit eBGP connections (i.e. TCP 179) only to/from expected peers. It would enhance security to permit only 'established' sessions from port 179.	NIST SP800-54 4.2.4
	BCP 38/84 + uRPF	Enforce source address validation for the data plane egressing towards peers where possible. Use uRPF (with ignore-default) at internal network boundaries where possible.	NIST SP800-54 4.2.5-6 CSRIC 5.1.3.4
	Record Peering State Changes (i.e. logging)	Peering state changes should be recorded for audit purposes in case they are indicative of security implications. Tools such as QUAGGA may be used for this purpose. Syslog of appropriate level towards a Security Incident and Event Monitoring (SIEM) function is most likely to be practical for this purpose.	NIST SP800-54 4.1
DOS / DDOS	CPU Overload	Mechanisms to prevent CPU/RP processing of irrelevant or malicious packets should be applied to limit the possibility of resource exhaustion attacks.	
	Counter DDoS	Service Providers should follow existing best practice with regards to BGP configuration and DDoS, as outlined in the suggested reference documents.	CSRIC 2A – Cyber Security Best Practices RFC 4778 BCP 46 NIST SP800-53

BGP Security		Standards and Best Practices	
Feature Type	Target Area	Recommendations	Key Reference Documents
General	BGP Route Refresh	To avoid associated memory use, RFC 2918 (Route Refresh) approach should be used where available.	NIST SP800-54 4.1
IPv6	De-aggregation prevention.	Given that the IPv6 address space size presents a very large opportunity for de-aggregation attacks (that could also break RIB scale), longest prefix recommendations must be adhered to and evolved over time. Recommendations shall be periodically reviewed in line with hardware evolution (e.g. larger route table scale). Base recommendation that prefixes longer than /48 are neither announced to nor accepted from the internet (with the /24 recommendation remaining for IPv4). Maximum prefix tracking per peer has increased significance due to potential de-aggregation scale.	RFC 7454
	IPv6 Special Purpose Prefixes	The IANA registry should be used in prefix-list configuration relating to special prefixes, with only 'Global: False' prefixes discarded for internet peerings.	RFC 7454
	Filtering	Given that IPv6 address space is only partially allocated, checks against IANA (and RIR) allocated address space in prefix-filters is recommended, with these updated on an iteration cycle of <1 month.	RFC 7454

References

1. Border Gateway Protocol (BGP) Good Practice Guide. CPNI, April 2011
2. BGP Operations and Security RFC7454. IETF Feb 2015
3. Mutually Agreed Norms for Routing Security (MANRS) 2014-2016
4. Border Gateway Protocol Security NIST 800-45 July 2007
5. BGP Security Best Practices – CSRIC March 2013
6. BGP prefix origin validation - https://tools.ietf.org/html/rfc6811
7. Resource Public Key Infrastructure - https://tools.ietf.org/html/rfc6480

Further reading:

8. https://www.manrs.org/resources/tutorials/irrs-rpki-peeringdb/
9. https://www.ripe.net/manage-ips-and-asns/resource-management/certification/resource-certification-roa-management
10. https://www.ripe.net/manage-ips-and-asns/resource-management/certification/resource-certification-rpki-for-provider-independent-end-users

This document is issued by the NCSC

For additional copies or general queries, please contact:

NCSC Enquiries

Hubble Road

Cheltenham

Gloucestershire

GL51 0EX

United Kingdom

Tel: 0300 020 0964

https://www.ncsc.gov.uk/contact

DISCLAIMER