TSA - CoP Technical Guidance Measures

Do telecom companies have to comply only with Telecoms Security Code of Practice (CoP) for now?

Are NCSC guidelines to be considered only as best practices?

Well, the CoP directly refers to some NCSC guidelines that are a minimum requirement

Below study reveals which CoP measures correspond to NCSC's minimum implementation standards

Telcos are advised to make serious efforts to comply with both i.e. CoP & NCSC guidelines

Analysis - Code of Practice Measures (Due by March 2024*) vs NCSC Guidelines

Measure Number	Description	Corresponding NCSC BGP Guideline No.	Corresponding NCSC BGP Guideline	Corresponding NCSC BGP Guideline / Minimum or Recommended	Techniques	Compliant
	Signalling Plane 1					
M3.11	External BGP updates shall be monitored for evidence of misuse.	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.	Minimum		
M3.12	Any BGP misuse that impacts a provider's network or services shall be mitigated in a timely manner, and at least within 12 hours whenever technically possible.					
M3.13	Providers shall ensure that contact details are current and accurate on all the Regional Internet Registries (e.g. RIPE) and should endeavour to keep other data sources accurate	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.	Minimum		
M3.14	All address space and autonomous system number (ASN) resources 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	12.2	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.	Minimum		
M3.15	Providers shall implement **ingress** and **egress** route filtering.	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	Minimum		

M3.16	Providers shall adopt and implement mechanisms that prevent IP address spoofing.	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	Minimum	Static Packet Filters. Manual Dynamic Packet Filters. Dynamically provisioned, RADIUS can be used Forwarding Based Validation. Using the router's forwarding table . Unicast Reverse Path Forwarding (uRPF) is the first of several technologies that provide BCP 38 capabilities. Network Address Translation (NAT). Inherent source validation	**Have you applied unicast RPF?** **Have you applied bogon list against BGP peers?**
M3.17	The provider shall share such details, as are appropriate and proportionate, of any BGP misuse with other providers where it may cause a connected security compromise.	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.	Minimum		
M3.18	An external path update that includes a prefix owned by the provider shall not be accepted.	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.	Minimum	13.1 Route Signing: Any LIR (Local Internet Registry) interested in protecting their network should sign their prefixes to accurately match their routing requirements. 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. Phase 1 - monitor mode Phase 2 - full filtering	**Have you implemented RPKI?** **Do you have AS_Path filters available to check invalid AS_Path?** **Do your BGP import/export policies filter prefixes and are they applied on peering policies?** **Have you applied BOGON filters?**
M3.19	End-users shall not be able to spoof IPs over the data plane (e.g. in line with BCP38).	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	Minimum	Unicast Reverse Path Forwarding (uRPF) is the first of several technologies that provide BCP 38 capabilities.	**Have you applied unicast RPF?**
	Third Party Supplier Measures 1					
M4.03	The provider shall record all equipment that remains in use but has reached the vendor's end-of-life date. Providers shall regularly review their use of this equipment, with a view to reducing the risk of a security compromise occurring as a result of unsupported equipment remaining in use.					**Are all the EOL/EOSupport devices listed in a repository?**
	Overarching security measures - Isolating Exposed Edge					
M1.05	Security boundaries shall exist between the **exposed edge** and critical or sensitive functions that implement protective measures.					
M1.06	Equipment in the **exposed edge** shall not be able to impact operation or routing within the core network. As an example, the exposed edge shall not be a PE-node within the provider's IP Core.					

Analysis - Code of Practice Measures (Due by March 2025) vs NCSC Guidelines

Measure Number	Description	Corresponding NCSC BGP Guideline No.	Corresponding NCSC BGP Guideline	Corresponding NCSC BGP Guideline / Minimum or Recommended	Techniques	Compliant
	Signalling plane 2					
M7.03	Providers shall have in place the means for recipients of their BGP routing updates to validate that the BGP routing update originated from the legitimate owner.	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.			
M7.04	Where the necessary information is available, providers shall validate that any BGP route updates they receive have originated from the legitimate owner.	**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.			

*	**March 2024 is the deadline for Tier 1 providers, for Tier 2 providers it is March 2025**

TSA CoP Compliance M1.03 to 1.06 – Isolating the Exposed Edge

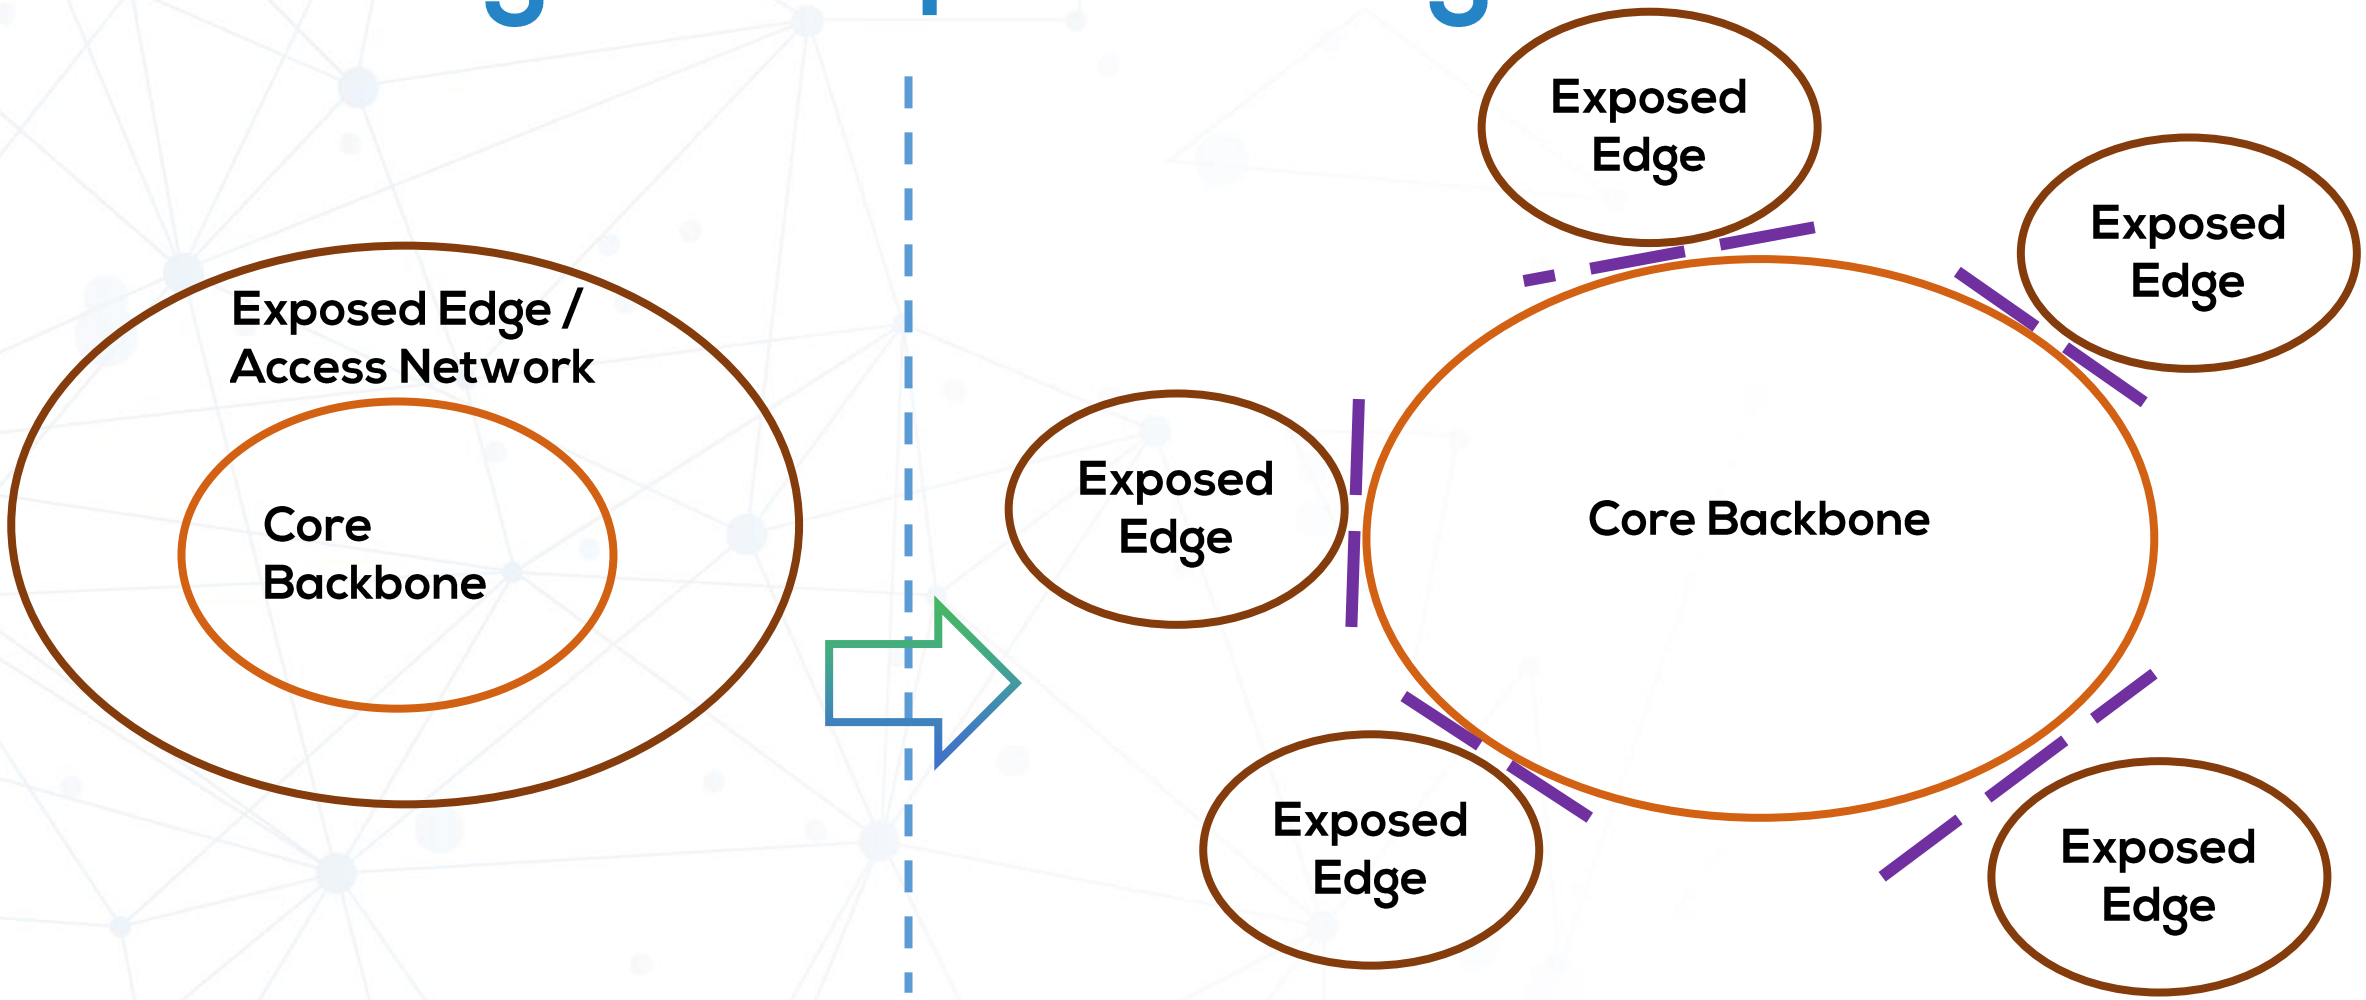

UURNIK SYSTEMS

Sample TSA/NCSC Compliance Report

Sr #	Feature Type	Target Area	NCSC recommendation	Peering	Recommendations/Advisories
1.0 Filtering					
1.1		Egress	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).	Yes	
1.2		Forwarding	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.	AS_Path filters not available to check invalid AS_Path	AS_Path filtering policies to be reviewed and amended according to guidelines
1.3		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.	No check for prefixes with lengthy AS_Path	AS_Path filtering policies to be reviewed and amended according to guidelines
2.0 Scalability					
2.1		Prefix-limits	Prefix-limits should be applied where expected prefix count can be known and managed for peers.	1000 prefixes	
2.2		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.	Not deployed, no routes pointing to null-0	Policies for blackholing undesired traffic are to be reviewed and designed

Sr #	Feature Type	Target Area	NCSC recommendation	Peering	Recommendations/Advisories
2.3		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.	Hardware filter present, generic system level policer present but it doesn't specify particular protocols thereby denying others, no CoPP for this peer	CoPP to be designed with rate-limiting of various different protocols
3.0 Stability					
3.1		Route-flap Dampening	Current recommendation is that it should not be used.	Not deployed, that doesn't let the router circumvent inadvertent behaviour due to route flaps	
3.2		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.	Not deployed, that stops the device from being able to continue packet forwarding throughout the control plane reset	Graceful restart to be implemented keeping in mind any implications on BFD
4.0 Security					
4.1		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).	Hash2	
4.2		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.	No ACL to control specific BGP peers	Specific BGP peers to be enlisted and controlled using ACLs

Sr #	Feature Type	Target Area	NCSC recommendation	Peering	Recommendations/Advisories
4.3		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.	No uRPF check available	
5.0 DOS/DDOS					
5.1		CPU Overload	Mechanisms to prevent CPU/RP processing of irrelevant or malicious packets should be applied to limit the possibility of resource exhaustion attacks.	No	
6.0 General					
6.1		BGP Route Refresh	To avoid associated memory use, RFC 2918 (Route Refresh) approach should be used where available.	Not deployed	