

DE-CIX INTERWIRE

GLOBEPEER TECHNICAL SERVICE DESCRIPTION

I. GENERAL PROVISIONS

1. Overview, scope of application

This document contains the Technical Service Description (TSD) for the GlobePEER product. This TSD is part of the DE-CIX INTERWIRE contractual framework for domestic internet access.

This TSD shall apply only to the GlobePEER product. The GlobePEER product may however be a prerequisite for other DE-CIX INTERWIRE services. This document contains only technical specifications and documentation. Please consult the GlobePEER PSSLA for service levels.

2. Amendment

This document may be revised and amended at any time pursuant to the provisions of the DE-CIX INTERWIRE INTERNET SERVICES PVT LTD (in the following called the DE-CIX INTERWIRE) Agreement.

3. Product prerequisites

The GlobePEER Product requires the following DE-CIX INTERWIRE products for its normal operation:

- DE-CIX INTERWIRE Access (see Master SLA and DE-CIX INTERWIRE Technical Access Description (TAD)) at any data centre location that allows a local connection to the DE-CIX INTERWIRE GlobePEER region.

4. Applicable standards

Members' use of the DE-CIX INTERWIRE network shall at all times conform to the relevant standards as laid out in [STD0001](#) and associated Internet STD documents.

II. DATA LINK-LAYER CONFIGURATION

1. Bandwidth

Bandwidth of the GlobePEER product must be explicitly configured, if the agreed bandwidth for GlobePEER differs from the bandwidth of the access or bundle of aggregated accesses, on which the GlobePEER product is used.

2. Frametypes

The following general policies shall apply:

<u>Frametype (ether types)</u>	<u>Policy</u>	<u>Enforcement</u>
0x0800 – IPv4 0x0806 – ARP 0x86dd – IPv6	Allow	-
All other types	Discard	Strict – all frames other than allowed types are dropped

3. MAC address configuration

All frames forwarded to the GlobePEER service shall have the same source MAC address.

4. Broadcast/Multicast Traffic

The following policies shall apply to broadcast/multicast traffic

<u>Protocol</u>	<u>Policy</u>	<u>Enforcement</u>
Broadcast ARP (excluding proxy ARP), multicast IPv6 Neighbor Discovery (ND)	Allowed, but rate limited to 1000kbps	-
All other types, i.e. including, but not limited to: - IRDP - ICMP redirects - IEEE802 Spanning Tree - Vendor proprietary discovery protocols (e.g. CDP) - Interior routing protocol broad/multicasts (e.g. OSPF, IS-IS, IGRP, EIGRP) - BOOTP/DHCP - PIM-SM - PIM-DM - DVMRP	Discard	Discarded, unless specifically allowed

III. IP LAYER CONFIGURATION (ISO/OSI LAYER 3)

1. Interface configuration

Interfaces connected to DE-CIX INTERWIRE ports shall only use IP addresses and netmasks (prefix lengths) assigned to them by DE-CIX INTERWIRE. The assignment will be provided in writing (e.g. email) during the provisioning process. In particular:

<u>Parameter</u>	<u>Policy</u>	<u>Remarks</u>
IP addresses (IPv4, IPv6) including subnet mask for your interfaces	IPv4 required	At least the IPv4 address has to be configured
IPv6 addresses (link-local & global scope)	No auto-configuration	All IPv6 addresses must be explicitly configured
IPv6 address (site-local)	Not allowed	IPv6 site-local addresses must not be used
Standard MTU	Fixed size	Standard IP MTU size must be explicitly set to 1500 Bytes, unless explicitly agreed in writing.

2. Routing configuration

The customer system's routing configuration shall include the following policies/settings:

<u>Parameter</u>	<u>Policy</u>	<u>Remarks</u>
BGP Version	v. 4 only	-
AS numbers	Public only	No AS numbers allowed from ranges reserved for private use across the entire DE-CIX INTERWIRE network.
Multiple ASN	Allow	Members may use more than one ASN for their DE-CIX INTERWIRE peering, provided that each ASN presented shares the same NOC and peering contact details.
Route advertising	Maximum aggregation	All routes advertised shall be aggregated as far as possible.
Route advertising – target IP	Advertising router only	All routes advertised across the Mumbai-IX network must point to the router advertising it unless an agreement has been made in advance in writing by Mumbai-IX and the members involved.
Route advertising – registration	Public registration required	All routes to be advertised in a peering session across Mumbai-IX must be registered in the RIPE database or another public routing registry.
IP-address space advertising	With permission only	IP address space assigned to Mumbai-IX peering LAN shall not be advertised to other networks without explicit permission of Mumbai-IX.
Mumbai-IX advertised routes	Accept	You can safely accept any routes announced by us, as all incoming advertisements are filtered according to the configured policies.

3. Route server feature

The Mumbai-IX route server system consists of two servers running BGP. For normal operation, only one is needed.

3.1 Minimum configuration

In order for the Mumbai-IX measurements of the route server feature to function, at least one connection to one route server must be set up with the following parameters:

<u>Parameter</u>	<u>Policy</u>	<u>Remarks</u>
connection mode	Active	Mumbai-IX side is configured as passive
bgp enforce-first-as	Not allowed	Enabled by default, must be disabled manually
AS-Set	Required	Mumbai-IX needs the customer AS-Set to build the filter rules
martians/bogons	Will be discarded	

3.2 BGP announcement validation

BGP announcement provided by the customer to the Mumbai-IX route server are validated for security reasons. For the validation route databases might be used (e.g. RADB).

3.3 Optional: communities

In addition to the one route server minimum configuration, the Customer may elect to control outgoing routing information directly on the Mumbai-IX route server by joining communities. Communities are processed by the Mumbai-IX route servers by the following set of filter rules:

<u>#</u>	<u>action</u>	<u>community</u>	<u>Local Preference</u>
1	block announcement of a route to a certain peer	0:<peer-as>	50
2	announcement of a route to a certain peer	<route-server-as>:<peer-as>	
3	block announcement of a route to all peers (monitoring only session)	0:<route-server-as>, no advertise, no-export	0
4	announcement of a route to all peers	<route-server-as>:<route-server-as> (default if nothing set)	100

Customers are kindly asked to consult the location specific documentation of existing communities, made available upon request.

4. Blackholing

Blackholing means diverting the flow of data to a different next hop (the “Blackhole”) where the traffic is discarded. The result is that no traffic reaches the original destination and hence hosts located within the "blackholed" prefix are protected from massive distributed denial of service (DDoS) attacks congesting the connection from the customer to Mumbai-IX. Thus blackholing is an effective way of mitigating the effects of DDoS attacks, etc.

Mumbai-IX provides the technical infrastructure to allow Blackholing to be set upped and used by customers. Mumbai-IX however have no control in cases where a customer is accepting these “Blackholed” prefixes.

4.1 Basic principle

4.1.1 In standard conditions

Customers advertise their prefixes with a Next Hop IP address belonging to their AS

- IPv4: /8 <= and <= /24

- IPv6: /19 <= and <= /48

4.1.2 In case of DDoS

Customers advertise their prefixes with a unique Mumbai-IX-provided Blackhole next hop IP address (BN)

- IPv4: /8 <= up to = /32 (if and only if the BN is set)
- IPv6: /19 <= up to = /128 (if and only if the BN is set)

Further, the standard announcement checks still apply.

4.2 L2 filtering

- Blackhole next hop (BN) has a unique MAC address (determined by ARP for the BN IP address) e.g. de:ad:be:ef:66:95
- ARP resolving for the Blackhole IP next hop is currently served by a host operated Mumbai-IX
- All edge nodes have a static entry for the unique MAC address
- Attack traffic is forwarded from the customer to the service with the static MAC address, traffic is denied.

4.3 Result

As a result, all traffic to the attacked and "blackholed" IP prefix is discarded already on the incoming switch, and hence victim's resources (e.g. connection from customer to Mumbai-IX) are protected.