

DE-CIX 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 contractual framework.

This TSD shall apply only to the GlobePEER product. The GlobePEER product may, however, be a prerequisite for other DE-CIX services. This document contains only technical specifications and documentation. Please consult the GlobePEER Special Service Level Agreement (Special SLA) for service levels.

2. Amendment

This document may be revised and amended at any time pursuant to the provisions of the DE-CIX Agreement.

3. Product prerequisites

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

- DE-CIX Access (see Master SLA and DE-CIX Technical Access Description (TAD)) at any data center location that allows a local or remote¹ connection to the respective GlobePEER region.

4. Applicable standards

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

¹ Some Exchange locations of DE-CIX are interconnected. At those locations customers can book the access to the GlobePEER region at the remote location as an additional service, e.g., customers of DE-CIX New York region can order the access to the DE-CIX GlobePEER Frankfurt region.

II. DATA LINK-LAYER CONFIGURATION (ISO/OSI LAYER 2)

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 access, on which the GlobePEER product is used.

2. Frame types

The following general policies shall apply:

<u>Frame type (ethertypes)</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) (Please note that Router Advertisement (ND-RA) packets are not included in this allowance)	Allowed, but rate limited to 1,000kbps	-
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 ports shall only use IP addresses and netmasks (prefix lengths) assigned to them by DE-CIX. 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
IP address of route servers	Required for credit claim	Configure at least one BGP session to one route server to be able to claim credits for the GlobePEER service. Advertising routes are not a requirement.

2. Additional configuration parameters

<u>Parameter</u>	<u>Policy</u>	<u>Remarks</u>
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 1,500 Bytes, unless explicitly agreed in writing.

3. 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 network.
Multiple ASN	Allow	Each GlobePEER service is defined by one ASN. If a customer wants to use multiple ASNs, multiple GlobePEER services must be used. Each ASN presented must have 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 DE-CIX network must point to the router advertising it, unless an agreement has been made in advance in writing by DE-CIX and the members involved.
Route advertising – registration	Public registration required	All routes to be advertised in a peering session across DE-CIX must be registered in the RIPE database or another public routing registry.
IP-address space advertising	With permission only	IP address space assigned to DE-CIX peering LAN shall not be advertised to other networks without explicit permission of DE-CIX.
DE-CIX advertised routes	Accept	You can safely accept any routes announced by us, as all incoming advertisements are filtered according to the configured policies.

4. Traffic forwarding

Traffic shall only be forwarded to a DE-CIX member, if permission has been given by the receiving member either:

- by advertising a route across the DE-CIX network (directly or via the route server)
- or explicitly in writing

5. Route server feature

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

5.1 Minimum configuration

In order for the DE-CIX measurements of the route server feature to function, and thus for a customer to be eligible for any credits, 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	DE-CIX Side is configured as passive
bgp enforce-first-as	Not allowed	Enabled by default, must be disabled manually
AS-Set	Required	DE-CIX needs the customer AS-Set to build the filter rules
martians/bogons	Will be discarded	

5.2 BGP announcement validation

BGP announcement provided by the customer to the DE-CIX route server are validated for security reasons. Databases might be used for the route validation (e.g. RADB).

5.3 Optional: communities

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

#	<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

The number and list of available communities may vary between GlobePEER regions and locations. Customers are kindly asked to consult the location-specific documentation of existing communities, made available upon request.

6. 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 DE-CIX. Thus blackholing is an effective way of mitigating the effects of DDoS attacks etc.

DE-CIX provides the technical infrastructure to allow Blackholing to be set up and used by customers. However, whether a certain customer accepts “Blackholed” prefixes or not is out of the control of DE-CIX.

6.1 Basic principle

6.1.1 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

6.1.2 In case of DDoS

Customers advertise their prefixes with a unique DE-CIX-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)

The standard announcement checks still apply.

6.2 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 the host buoy.
- 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 ingress. This results in attack traffic not leaving the node through which it enters the GlobePEER service and it is discarded locally.

6.3 Result

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