

FlexTunnel

KNOWLEDGE BASE - OCTOBER 2026

IN THIS GUIDE

01 Service overview
02 Service boundary and traffic direction
03 GRE and packet handling
04 Overlay addressing and route exchange
05 Bandwidth and traffic treatment
06 Burstable-service measurement
07 Multiple tunnels and path selection
08 Optional routing controls
09 Responsibilities
10 Service levels and support

> This document is an informative description of the standard Inter.link FlexTunnel service. The applicable service order, Service Level Agreement (SLA), and Master Services Agreement (MSA) prevail in the event of any conflict.

01

Service overview

Inter.link FlexTunnel is a standalone GRE-based Layer 3 service delivered over the public IPv4 Internet. It allows Inter.link to deliver unicast IPv4 and IPv6 traffic to a customer without requiring a physical Inter.link port or cross-connect at the customer location. The customer provides a publicly reachable IPv4 tunnel source, terminates GRE, and establishes BGP with Inter.link AS5405 over the tunnel.

FlexTunnel can terminate only on selected Inter.link routers. Current availability must be confirmed in the customer portal or with Sales. The service order records the customer-specific endpoints, bandwidth, term, and charges.

FlexTunnel may be used independently or as the clean-traffic return path for a separately ordered DDoS Protection service. It is not an Internet transit route feed, a Layer 2 service, or an encrypted VPN. GRE is the only supported tunnel type; IPsec and other encrypted variants are not available.

02

Service boundary and traffic direction

BGP control traffic is exchanged bidirectionally over the GRE tunnel. The customer data plane is unidirectional: Inter.link delivers traffic toward customer prefixes, while customer-originated payload traffic toward Inter.link is outside the service definition. The customer must not point a static route or default route toward Inter.link over FlexTunnel.

Only unicast IPv4 and IPv6 payloads are supported. Multicast, non-IP Ethernet payloads, and general Internet transit toward Inter.link are not supported. Inter.link does not advertise unicast routes to the customer over FlexTunnel.

03

GRE and packet handling

FlexTunnel uses the GRE base header without optional checksum, key, or sequence number fields.

PARAMETER	STANDARD BEHAVIOUR
Underlay	Public IPv4 Internet
Customer tunnel source	One customer-supplied, publicly reachable IPv4 address per tunnel
Inter.link tunnel destination	Assigned during provisioning
GRE key, checksum, and sequence number	Not used
Outer IPv4 TTL	Fixed at 255
Outer IPv4 DF	Set; outer fragmentation is not supported
Path MTU discovery	Enabled on the Inter.link GRE endpoint
Supported inner payload	Unicast IPv4 and IPv6
Inner IP MTU	Static 1,476 bytes for both IPv4 and IPv6
TCP MSS	Customer-side clamping is required; recommended maxima are 1,436 bytes for IPv4 and 1,416 bytes for IPv6, with lower values where additional headers require them

The inner DSCP and ECN fields are preserved. The outer IPv4 DSCP field is set to CS0 and the outer ECN field to Not-ECT; neither is copied from the inner packet. No quality-of-service treatment is guaranteed across the public underlay.

Inter.link does not fragment oversized inner or outer packets. Traffic that cannot be delivered without fragmentation is dropped, and customers must assume that no ICMP error will be returned. The inner MTU is not dynamically reduced when a lower underlay path MTU is discovered.

04

Overlay addressing and route exchange

Inter.link assigns an IPv4 `/31` and an IPv6 `/127` overlay transfer network per tunnel by default. One IPv4 and one IPv6 BGP session are provided, and each session carries matching-family customer unicast routes. Additional sessions are supported. Both public and private customer ASNs may be used.

Session parameters

PARAMETER	STANDARD BEHAVIOUR
Provider ASN	AS5405
Protocol and transport	BGP-4 over directly connected sessions within GRE
BGP keepalive	60 seconds
BGP hold time	180 seconds
TCP MD5	Supported and encouraged, but not required
GTSM	Supported, disabled by default
BFD	Optional asynchronous BFD for IPv4 and IPv6; transmit 300 ms, receive 300 ms, multiplier 3
Graceful restart	Enabled with restart time 300 seconds

Inter.link acts as a graceful-restart helper and retains customer routes when the customer advertises graceful-restart capability for the relevant IPv4 or IPv6 address family. During an Inter.link control-plane restart, forwarding is preserved and Inter.link advertises the Forwarding State bit. A BFD failure or confirmed tunnel-path failure overrides graceful restart and causes immediate route removal.

Customer advertisements use the same route filtering and routing policy as IP Transit. Detailed route validation, route acceptance, traffic-engineering communities, and RTBH policy are defined in the separate Routing Policies document.

Maximum-prefix management

Maximum-prefix limits are derived from the customer's PeeringDB record and synchronized daily. Missing PeeringDB information is verified manually. Stale information leaves the configured limit unchanged, while a reduced PeeringDB value is applied at the next synchronization. Automatic synchronization may be disabled on request; subsequent manual limit changes are requested by opening a ticket in the customer portal.

Exceeding the limit does not shut down the BGP session. Additional prefixes are ignored. The portal displays BGP session status, the current limit, and accepted and ignored prefix counts.

05

Bandwidth and traffic treatment

The ordered CDR is a billing commitment and not a data-plane policer or shaper. Traffic may burst above it, with no configured service ceiling below the available end-to-end capacity. Current CDR options and prices must be confirmed in the customer portal or with Sales.

Packet forwarding is best effort and subject to acceptable use.

06

Burstable-service measurement

Burstable usage is measured using inner IP bytes; outer Ethernet, IPv4, and GRE overhead is excluded. For each target month, Inter.link obtains inbound and outbound traffic data in five-minute intervals and calculates a 95th-percentile value independently for each direction. The greater value is the measured monthly usage.

For a separately ordered aggregation group, the independently selected p95 value of each member service is summed and compared with the aggregation-group CDR. Aggregation may include services at different points of presence (PoPs).

07

Multiple tunnels and path selection

One or more tunnels may be ordered, and no redundant tunnel is required. Every tunnel is forwarding-ready, with traffic selection governed by normal BGP best-path selection. Bandwidth may be ordered independently per tunnel or included in a separately ordered billing aggregation group.

Tunnels may terminate on different Inter.link routers, but different-router termination does not guarantee diversity of the public underlay, routing, or other path components.

08

Optional routing controls

RTBH is included and follows the community and route-acceptance rules in the separate Routing Policies document. It does not provide managed detection, traffic scrubbing, or clean-traffic delivery.

BGP FlowSpec is a separately ordered, customer-wide add-on that can be used across the customer's IP Transit and FlexTunnel services. IPv4 rules use the existing IPv4 session and IPv6 rules use the existing IPv6 session. A customer using FlowSpec over FlexTunnel requires at least 10 Gbps of aggregate FlexTunnel commitment; commitments from multiple FlexTunnel services may be combined.

DDoS Protection is also ordered separately. When FlexTunnel is used for clean return, the tunnel carries only customer-bound clean traffic and Inter.link does not advertise routes over it. Clean return has no separate technical ceiling beyond the available end-to-end capacity; the ordered CDR and burst billing continue to apply.

09

Responsibilities

ACTIVITY	INTER.LINK	CUSTOMER OR THIRD PARTY
Tunnel endpoint	Operates the Inter.link endpoint and assigns the overlay addresses	Supplies a publicly reachable IPv4 source and operates the customer endpoint
Underlay	Provides the Inter.link-side Internet connectivity	Provides customer-side Internet connectivity and permits GRE through intervening firewalls
Routing	Establishes its side of BGP and applies routing policy	Configures BGP, maintains route authorizations, and does not direct customer-originated payload traffic toward Inter.link
MTU and packet handling	Sets the fixed inner MTU and outer-header behaviour	Configures the 1,476-byte inner MTU and required TCP MSS clamping
Capacity	Measures inner IP bytes and delivers traffic over the available path	Selects CDR or aggregation and manages customer-side capacity
Multiple tunnels	Keeps each ordered tunnel forwarding-ready	Controls path selection through BGP and does not assume different-router termination provides path diversity
Security	Operates the service endpoint; GRE itself provides no encryption	Protects the endpoint and uses application or network-layer encryption where required
Fault isolation	Manages the Inter.link endpoint and support ticket	Provides underlay evidence and customer-side testing

10

Service levels and support

Availability and performance commitments, incident priorities, response and restoration targets, maintenance rules, exclusions, remedies, and service credits are defined exclusively in the applicable SLA. Nothing in this informative specification changes those terms.

Need a hand?

Our network engineers are happy to help — reach out through the Inter.link Portal or visit inter.link.