

FlexEthernet

KNOWLEDGE BASE - OCTOBER 2026

IN THIS GUIDE

01 Service overview
02 Physical and logical handoff
03 Layer 2 transparency
04 MAC learning and forwarding
05 Bandwidth and traffic treatment
06 Resilience and diversity
07 FlexConnect
08 Responsibilities
09 Service levels and support

> This document is an informative description of the standard Inter.link FlexEthernet 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 FlexEthernet is a private Layer 2 Ethernet transport service. It connects two or more customer endpoints in a single Ethernet broadcast domain across the Inter.link network. Point-to-point services have two endpoints, while multipoint services have three or more.

Endpoints may use different physical port speeds, VLAN presentations, and ordered bandwidths. Multiple endpoints may terminate on the same physical port or Link Aggregation Group (LAG). Adding or removing an endpoint preserves traffic between the remaining endpoints.

The service is available at every Inter.link point of presence (PoP). The service order records the endpoints, ports, bandwidths, VLAN presentations, diversity, term, and charges. Current port speeds and bandwidth options must be confirmed in the customer portal or with Sales.

The standard service boundary extends from the Inter.link provider edge to the Inter.link patch panel in each building meet-me room (MMR). Customer equipment, cabling, power, and any extension beyond that demarcation remain the customer's responsibility. A data-centre cross-connect may be ordered directly by the customer or supplied by Inter.link as a separately quoted item.

02

Physical and logical handoff

All physical interfaces operate in full-duplex mode. Auto-negotiation is not supported. Inter.link supplies and maintains the transceiver installed in Inter.link equipment; the customer supplies and maintains the transceiver in the customer equipment.

Supported transceivers

INTERFACE	STANDARD TRANSCEIVER	FEC BEHAVIOUR
1G	1G-LX	—
10G	10G-LR	—
25G	25G-LR	RS-FEC by default; FC-FEC is available
100G	100G-LR or 100G-LR4	RS-FEC where required by the optic
400G	400G-LR4	—

All listed optics use single-mode fibre. Connector type, reach, wavelength, and optical operating limits follow the selected transceiver. Other optical options may be quoted separately.

Ethernet presentation

Each endpoint may use a tagged or untagged handoff. Only one untagged service may be configured on a physical port. For a tagged handoff, the outermost VLAN ID must exactly match the configured service VLAN. Inter.link translates the outermost VLAN at egress, allowing each endpoint to use a different VLAN ID.

ATTRIBUTE	STANDARD BEHAVIOUR
Tagged service	Outer TPID `0x8100`; VLAN IDs 1–4094 are available per port or LAG
Untagged service	One untagged service per physical port
Additional tags	Tags below the outermost service tag are preserved; there is no explicit tag-depth limit within the service MRU
PCP and DEI	Preserved but not interpreted or honoured for traffic priority
Ethernet MRU	9,100 bytes for the complete frame as received, including Ethernet headers, VLAN tags, payload, and FCS
IEEE 802.3x PAUSE and PFC	Transported transparently as customer Layer 2 traffic; not used to provide a service class or flow-control guarantee

Link-loss propagation is not provided. Loss of an endpoint link is not propagated to the other service endpoints.

Link aggregation

LACP is supported with one or multiple members. All members must have identical speeds and terminate on the same Inter.link router. Fast LACP timers are used, and no minimum members are configured.

Traffic is distributed using available Layer 2, IPv4, IPv6, and TCP or UDP header fields. Hashing preserves flows under normal operation but does not guarantee equal byte distribution between members. If a member fails, flows are redistributed across the remaining active members while the LAG remains operational.

03

Layer 2 transparency

FlexEthernet transports otherwise valid Ethernet frames within the service MRU without interpreting their EtherType. The following matrix records observed transparency tests. A passed test means that the traffic was observed to pass under the tested conditions; it is not an unconditional guarantee for every packet variant or future implementation. Not tested means that no conclusive result is asserted and does not indicate a known failure.

CONTROL PROTOCOL OR TRAFFIC	UNTAGGED	SINGLE TAGGED	DOUBLE TAGGED
ARP	Passed	Passed	Passed
BFD Control	Passed	Passed	Passed
BFD Control (Micro-BFD)	Passed	Passed	Passed
BFD Echo	Passed	Passed	Passed
BFD Echo (Async)	Passed	Passed	Passed
BGP	Passed	Passed	Passed
CDP	Passed	Passed	Passed
CFM	Passed	Passed	Passed
EIGRP v2	Passed	Passed	Passed
E-LMI	Not tested	Not tested	Not tested
LACP	Passed	Passed	Passed
LLDP	Passed	Passed	Passed
ICMP	Passed	Passed	Passed
ICMPv6	Passed	Passed	Passed
IGMPv3 Query	Passed	Passed	Passed
IGMPv3 Response	Passed	Passed	Passed
IS-IS L1 Hello	Passed	Passed	Passed
IS-IS L2 Hello	Passed	Passed	Passed
ES-IS IS-to-ES	Not tested	Not tested	Not tested
ES-IS ES-to-IS	Not tested	Not tested	Not tested
IPv4 Multicast	Passed	Passed	Passed
IPv6 Multicast	Passed	Passed	Passed

CONTROL PROTOCOL OR TRAFFIC	UNTAGGED	SINGLE TAGGED	DOUBLE TAGGED
MACsec	Passed	Passed	Not tested
MLDv2 Query	Passed	Passed	Passed
MLDv2 Response	Passed	Passed	Passed
OAM	Not tested	Not tested	Not tested
OSPFv2 Hello	Passed	Passed	Passed
OSPFv2 LSA	Passed	Passed	Passed
OSPFv3 Hello	Passed	Passed	Passed
OSPFv3 LSA	Passed	Passed	Passed
PVST+	Passed	Passed	Passed
RARP	Passed	Passed	Passed
STP/RSTP/MSTP	Passed	Passed	Passed
UDLD	Not tested	Not tested	Not tested
VRRP v2	Passed	Passed	Passed
VRRP v3	Passed	Passed	Passed
VTP	Not tested	Not tested	Not tested

Inter.link does not terminate or operate customer CFM, Y.1731, Ethernet OAM, or loopback functions; those frames are carried as customer traffic. LACP transported as customer payload is distinct from LACP terminated on an Inter.link access LAG. Encryption, including MACsec, is entirely customer-managed.

04

MAC learning and forwarding

FlexEthernet dynamically learns source MAC addresses. A service supports 1,000 learned MAC addresses across all endpoints, independent of endpoint count. Entries age after 300 seconds of inactivity. Once the learning limit is reached, new addresses are not learned, but traffic continues to be forwarded.

A MAC address that moves between endpoints five times within 180 seconds has any further moves suppressed. Broadcast, unknown-unicast, and multicast traffic is replicated to every endpoint except the ingress endpoint. Endpoint isolation, storm control, and loop protection are not provided; the customer is responsible for preventing Layer 2 loops.

05

Bandwidth and traffic treatment

Bandwidth is ordered independently per endpoint and is symmetric. Traffic entering the Inter.link service is policed to the endpoint's ordered bandwidth. The ingress policer permits bursts within its configured burst allowance; traffic exceeding the policer profile is discarded. Destination egress is not policed, and contention is handled on a best-effort basis without per-sender fairness.

The policer measures the complete Ethernet frame, including headers, VLAN tags, payload, and FCS. Customer priority information is preserved but is not used to provide class-based QoS.

06

Resilience and diversity

FlexEthernet may be ordered with varying levels of diversity. Sales can confirm which diversity levels are available.

Each path receives its full ordered bandwidth and remains forwarding-ready. Because FlexEthernet is a Layer 2 service, the customer controls path selection and is responsible for preventing loops when multiple paths are connected.

07

FlexConnect

A FlexConnect Key allows another party to connect to a specific port and VLAN. The key is generated and may be revoked in the customer portal. It is single-use and does not expire. The party that uses the key to create the connection is billed; the party that generated the key is not billed, and the charge is not split.

08

Responsibilities

ACTIVITY	INTER.LINK	CUSTOMER OR THIRD PARTY
Network and provider edge	Operates the provider edge and transport network	—
Service configuration	Provisions the ordered endpoints, VLAN translations, and ingress bandwidth	Supplies complete and accurate order data
Transceivers and customer equipment	Supplies and maintains the Inter.link-side transceiver	Supplies and maintains the customer-side transceiver, equipment, cabling, power, and configuration
Cross-connect	May supply it when separately quoted	Orders and maintains it unless included in the Inter.link quote
Capacity	Delivers the ordered endpoint bandwidth	Selects sufficient port capacity and manages customer-side contention
Layer 2 design	Provides the ordered broadcast domain	Prevents loops and manages any customer-controlled encryption
Diverse delivery	Provides full ordered bandwidth on each path and the ordered diversity boundary	Operates all customer-side paths and controls forwarding topology
Fault isolation	Tests to the demarcation and manages the support ticket	Provides access, evidence, and customer-side testing

09

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.