



CD001

**Service Provider
Technical Overview Document**

52 Main Road, Hope 7020
PO Box 3005
Richmond 7050
Nelson, New Zealand
Fax: + 64 3 989 3631
Email: info@thelink.net.nz
Website: www.thelink.net.nz

Notice

All information contained herein is proprietary to **thelink** and no portion may be reproduced, stored in a retrieval system, transmitted in any form by any means, without the prior written approval of **thelink**. This document will be held in strict confidence by the recipient and will not be used, in whole or in part, for any other purpose other than the purpose for which it is provided without the prior written consent of **thelink**. In no event shall **thelink** be liable to anyone for any damages arising out of the use of this document.

© **thelink** 2009

Contents

1.	ABSTRACT.....	4
2.	TERMINOLOGY.....	4
3.	NETWORK SERVICE	5
4.	THELINK™ ETHERNET-SPA PHYSICAL INTERFACE.....	5
5.	THELINK™ ETHERNET-SPA LOGICAL INTERFACE.....	6
6.	THELINK™ ETHERNET-CA PHYSICAL INTERFACE	6
7.	THELINK™ ETHERNET-CA LOGICAL INTERFACE.....	6
8.	BANDWIDTH MANAGEMENT	7
9.	QUALITY OF SERVICE.....	7
10.	MULTICAST SUPPORT	7
11.	MAXIMUM FRAME SIZE	8
12.	NETWORK TRANSPARENCY	9
13.	VERSION CONTROL	9

1. Abstract

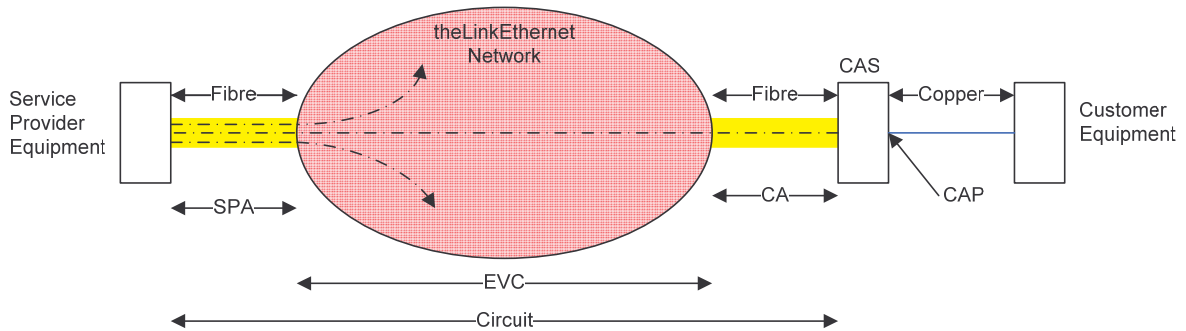
This document has been prepared to provide an overview of the **thelink**TM Ethernet technical specification from a Service Providers perspective.

2. Terminology

C-Tag	Customer VLAN/priority tag. The C-Tag, if present, is carried transparently through the network.
Circuit	A point-to-point path through an EVC network from an SPA to a CA.
CIR (Committed Information Rate)	Bandwidth usage through an EVC below which traffic policing is not applied.
Customer	Customer of the Service Provider for whom circuits are provided.
CA (Customer Access)	A fibre pair from the network to the customer premises, terminated on an CAS.
CAP (Customer Access Port)	A port on an CAS.
CAS (Customer Access Switch)	A small switch located at customer premises terminating a CA and providing a number of CAPs.
Double Tagged Packet	A Packet with two 802.1Q tags (both VIDs non-0), the first is the S-Tag, the second the C-Tag.
EVC (Ethernet Virtual Circuit)	A virtual path through the core network between an SPA and CA. Bandwidth is allocated on a per-EVC basis.
PCP (Priority Code Point)	3-bit priority field defined by 802.1P.
PIR (Peak Information Rate)	Maximum allowable bandwidth through an EVC, above which traffic policing is unconditionally applied. This is also known as EIR under the terms specified by the Metro Ethernet Forum.
S-Tag	Service VLAN/priority tag. The VID field of the S-Tag selects the service to which the packet belongs. The VID must not be set to 0.
SP (Service Provider)	A provider of network services leasing circuits through the network to Customers.
SPA (Service Provider Access)	A port connected directly to the network over which multiple circuits are provisioned.
Tagged Packet	Packet with an 802.1Q tag with VID field not set to 0.
Untagged Packet	Packet with no 802.1Q/P tag.

3. Network Service

thelink™ Ethernet service consists of circuits from an aggregated Service Provider Access (SPA) points, through Ethernet Virtual Circuits within the network to a Customer Access Port (CAP).



The service corresponds to the Metro Ethernet Form "E-Line" specifications. Circuits have exactly two endpoints, and have CIR/PIR determining the bandwidth available to that circuit.

Note that a service provider may oversubscribe the SPA as required.

4. **thelink™** Ethernet-SPA Physical Interface

Service provider access (SPA) is provided to the service provider as a tagged 1 Gbps or 10 Gbps Ethernet, direct to the service provider's port. Circuits on the SPA are delineated using 802.1Q VLAN tags.

The SPA interface is delivered as a lit fibre service over either a single fibre or a fibre pair depending on the service required. The interface is delivered to the provider's premises and terminated on SC connectors. The provider is responsible for obtaining a suitable optical module and transition cable.

For 1 Gbps access, 1000Base-BXU (TX1310 nm; RX 1550) single fibre optics are required. For 10 Gbps access, 10GBase-LR (1310 nm) optics are used.

Where the SPA fibre length exceeds 10km, alternative optics may be required. This will be specified prior to order acceptance, and may incur an additional installation charge.

5. thelink™ Ethernet-SPA Logical Interface

At the SPA, circuits are delineated using 802.1Q VLAN tags. Untagged traffic will be discarded by the network.

Service providers assign 802.1Q service VLAN IDs (S-VIDs) as part of the service order. S-VIDs should be in the range 2 through 3999; S-VIDs outside this range will not be accepted. (S-VIDs above 4000 are reserved for common services and future use. 1 is commonly used by equipment as a default VLAN number, and its use can cause compatibility and configuration issues.)

Double tagging (Q-in-Q) is supported on the SPA. Inner tags (C-tags) are passed transparently through circuits, and the network supports the extended frame sizes required for this.

802.3X flow control is not supported on the SPA.

6. thelink™ Ethernet-CA Physical Interface

Customer Access (CA) is provided via a Customer Access Switch (CAS) located at the customer's premises, connected to the network via a single fibre strand operating at 1 Gbps. The CAS provides a number of Customer Access Ports (CAPs).

The CAP interface to the customer premises equipment (CPE) is a copper Ethernet interface on the CAS, auto-negotiating to 1000BASE-T, 100BASE-TX or 10BASE-T, full or half duplex.

Each configured circuit is assigned to a separate CAP.

Note that only one CAS is provided per customer site, unless the sum of the CIRs of the circuits provisioned to the site exceeds 1 Gbps. Note that where multiple service providers have circuits provisioned to the same customer, they may be provided on the same CA and CAS.

The Service Provider is responsible for ensuring that the customer has a suitable interface and cabling to reach the CAS.

7. thelink™ Ethernet-CA Logical Interface

The CA maps each CAP to a separate circuit. Thus no special signalling is required. The CAS provides sufficient separation that circuits provided on two CAPs can be bridged together by customer equipment.

802.1P and 802.1Q tags are passed transparently through the network. 802.1P tag values are interpreted at congestion points in the network; see Quality of Service below.

802.3X flow control can be negotiated on the CAP. However, note that flow control information is not passed beyond the CAS.

8. Bandwidth Management

Multiple bandwidth CIR/PIR steps are provided:

Customer Access Circuits

- 10 Mbps PIR, 10 Mbps CIR
- 30 Mbps PIR, 30 Mbps CIR
- 100 Mbps PIR, 100 Mbps CIR
- 1000 Mbps PIR, 300 Mbps CIR

Service Provider Access Circuits

- 50 Mbps PIR, 50 Mbps CIR
- 100 Mbps PIR, 100 Mbps CIR
- 500 Mbps PIR, 500 Mbps CIR
- 1000 Mbps PIR, 1000 Mbps CIR

The network is designed not to congest provided that traffic rates remain below CIR. Above CIR, traffic may burst up to PIR, but may experience congestion.

Note that congestion may occur at the SPA due to over-subscription of the SPA. Such congestion does not represent a violation of CIR/PIR parameters.

9. Quality of Service

The network provides an uncontended virtual leased line service within the specified PIR/CIR steps, and traffic is forwarded through the network without re-ordering. Customers wishing to apply QoS to a circuit should implement queuing at or below the specified CIR rate limits on the ports facing the network as appropriate to their application.

Note that congestion may occur at the SPA due to over-subscription of the SPA. Such congestion does not represent a violation of CIR/PIR parameters.

10. Multicast Support

As circuits are point-to-point, multicast and broadcast packets are forwarded through from ingress to egress like unicast packets. The network does not attempt to prune multicast traffic. Multicast pruning should therefore be done in provider and/or customer equipment.

11. Maximum Frame Size

The network supports "jumbo" Ethernet payloads (MTU) of up to 9,000 bytes on interfaces negotiating link speeds of 1 Gbps or more.

SPA interfaces support frame sizes of up to 9,022 bytes, allowing for a layer 3 payload of 9,000 bytes, plus 14 byte MAC header and up to two 4-byte 802.1Q tags.

CAPs operating at 1 Gbps support frame sizes of up to 9,018 bytes, allowing for a 9,000 byte payload, 14 byte Ethernet header and optional 4-byte 802.1Q tag.

For CAPs negotiating to 10 or 100 Mbps, jumbo support is not available, and the maximum frame size is limited to 1518 bytes. Service providers should ensure that customer premises equipment both supports jumbo frames and can reliably negotiate 1 Gbps access with the CAS before enabling jumbo frame support.

12. Network Transparency

For the most part, the network carries all unicast and multicast traffic transparently, only processing 802.1Q S-tags as required. However, the Ethernet standards reserve a range of multicast MAC addresses for link protocols, which are normally not passed through bridge-type devices (including switches). This address range is 01-80-c2-00-00-00 through 01-80-c2-00-00-0F, and is used by several link-level protocols, including:

- 802.1D Spanning Tree, 802.1W Rapid Spanning Tree, MSTP etc
- 802.3X Flow Control
- 802.1AX Link Aggregation Control Protocol (LACP)
- 802.1X Access Control
- 802.1AB Link Layer Discovery Protocol (LLDP)

Of these, the following protocols are passed transparently through **thelink**TM Ethernet virtual circuits:

- Spanning Tree family protocols, including RSTP and MSTP (multicast address 01-80-C2-00-00-00);
- 802.1X Access Control (01-80-C2-00-00-03);
- 802.1AB Link Layer Discovery Protocol (01-80-C2-00-00-0E);

Other protocols in the link layer range may be passed, interpreted within the network or dropped. In particular, 802.3X flow control (01-80-C2-00-00-01) and 802.1AX LACP (01-80-C2-00-00-02) will not be passed.

Link-layer protocols outside the reserved range, including GVRP/MVRP, CDP and VRRP are passed transparently.

13. Version Control

Version Number	Change	Author	Date
V1.0	Draft Document for Review	Don Stokes	21 August 2009
V1.1	Review for Branding	Jodi Lively	December 2009
V1.2	Section 4 altered	Andrew Stanton	7 January 2010
V1.3	Additional bandwidth products	Andrew Stanton	24 March 2010
V1.4	Additional bandwidth products	Andrew Stanton	22 July 2011