



Application Note

ML-IP Networks for ISPs

Configure and deploy ML-IP for increased bandwidth and reliability

CONTENTS

1	Introduction	2
2	Network Design	3
2.1	Overall design	3
2.2	Design FAQs	4
3	Installation and Configuration	6
3.1	Server Side Unit	6
3.2	Client Side Unit	7
4	Conclusion	8

1 Introduction

This application note deals with setting up typical, simple infrastructure to allow an ISP to provide ML-IP enabled broadband services. It is mainly concerned with the network design, but deals briefly with installation and configuration, and common issues encountered during this process.

The design of a typical ML-IP ISP network is fairly simple, but does require understanding of the principles at work. ML-IP networks use tunnels (like VPN tunnels, although without encryption) connecting ePipe 2344s and ML-IP Concentrators. By controlling what links are used to transport data over the tunnel, the network's speed and reliability can be increased.

Deploying ML-IP in your network allows you to:

- Provide scalable bandwidth to customers
- Serve otherwise inaccessible or impractically priced locations, by:
 - Increasing network bandwidth.
 - Increasing network reliability.
- Sell network services which would otherwise be unavailable
 - Double, treble or quadruple the upstream and downstream speed of Internet connections at a linear increase in cost.
 - Fractional fiber bandwidth over standard DSL circuits.
 - VDSL-speed broadband services for multi-tenant, campus and Internet café installations, using multiple high speed links

DRAFT

2 Network Design

2.1 Overall design

As ML-IP creates tunnels, the routing and IP addressing for an ML-IP network is quite simple – it can be seen as another network, reachable through the ML-IP Concentrator which is installed at the ISP. The tunnel traffic itself is simply sent to the links (ADSL, PSTN, or whatever is in use) as any other traffic being sent to the subscriber.

A typical network design will look like this (simplified):




Figure 1

In this case, an ePipe 2344 ML-IP Gateway is installed at the client site, connected to up to three ADSL links. Note that you could use an ML-IP Concentrator at the client site if it is required – see chapter 2.2 for more information on which model to use under which circumstances.

The 2344 is the client end of the ML-IP tunnel. Note that most tunneling protocols do not typically use the client/server terminology, however ML-IP does. The client end of the tunnel initiates the connection, but once it is connected the two ends are essentially the same and data flows in both directions. In this example, the tunnel will be the default route for the 2344. Note that only a client side ML-IP tunnel can be the default route for an ePipe, server side tunnels must route to specific addresses or address ranges.

The placement of the ML-IP Concentrator (one Concentrator can serve up to 100 clients, each with multiple links) in an ISP network is fairly simple – it must be reachable by the client side links over IP. Typically the unit is configured as a 'one-armed ePipe' – that is, it has only one Ethernet interface. ML-IP packets flow into this connection, are un-encapsulated, and the original data packet goes out the same connection (this is reversed for incoming data of course).

This means that data going to the client site (the network behind the 2344) needs to be routed to the ML-IP Concentrator. The Concentrator will then send the data, encapsulated in ML-IP packets, to the links connected to the 2344, where the data is un-encapsulated and passed to the LAN as a single stream of data.

2.2 Design FAQs

Q. Should I use a 2344 or an ML-IP Concentrator at the client site? What are the differences between the two?

A. You can use either – in certain circumstances, either product may be the best choice. The 2344 is a small, appliance-like router, the ML-IP Concentrator is software which runs on a Linux PC (which could be a server, a small footprint PC, or an Intel based appliance).

- 2344 is a small, neat desktop unit and has no moving parts – this is the main reason for using one.
- Concentrator allows higher throughput, allows for higher number of incoming tunnels, and more interfaces.
- The Concentrator also allows for running other applications (proxy, web or mail servers for instance) on the Linux box, but this is not necessarily recommended.
- If only one interface is being used, we recommend using the Concentrator.
- If total required throughput is more than 6mbps, use the Concentrator.
- The Concentrator can be installed on an 'appliance-style' PC for these types of circumstances. We call this a 'distracted Concentrator'. An application note dealing with this topic available from our website.
- If using the Concentrator as a CPE device, up to four Ethernet NICs can be used as Internet connections, one more than the 2344 allows.
- Think of the Concentrator as an ePipe, which just happens to use a PC's hardware.

Q. Where should the server side device be installed?

A. Anywhere with sufficient bandwidth, and IP on Ethernet connectivity to the links that are being used at the client site. This depends on the particular ISPs architecture – it may be best to put it in near the core router, or to locate it next to web or mail servers, for instance, if the former is not possible.

Q. Why use only one Ethernet interface on the server side unit?

A. It simplifies routing – the server side unit needs to be able have a default route to contact the Internet, and also be able to contact the links on the client side units. If we were to use two ports (which must be on different IP subnets) and route between them, it would be necessary to set up static routes for all the client side units, as well as provide another IP subnet for the port to use. Using a single port avoids these complications.

Q. What IP addressing needs to be considered?

A. Addressing for four IP networks needs to be considered during the design phase of setting up an ML-IP network – however only simple checks are required for all of them. The networks are shown in Figure 2 below.

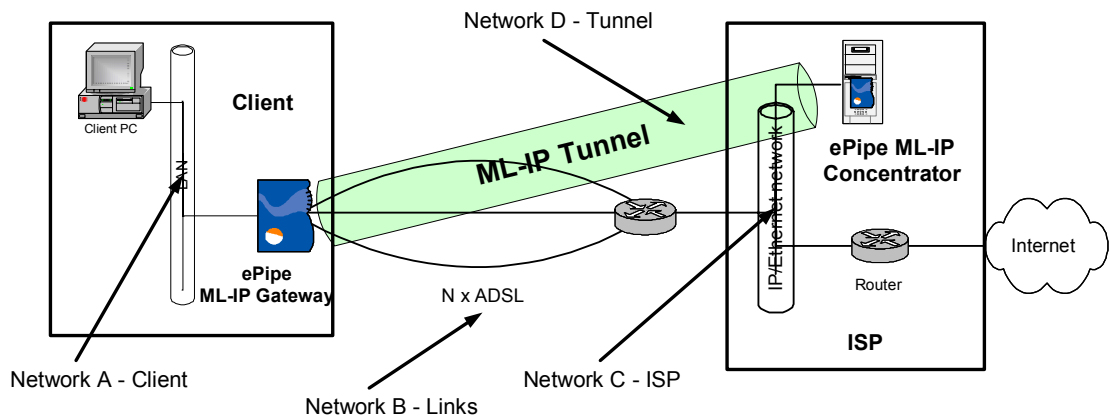


Figure 2

- Network A is the IP network immediately behind the ML-IP device at the client site. It must have public IP addresses, and these IP addresses must be routed from the Internet to the server side ML-IP device (in future ML-IP software versions, NAT will be available on traffic going across the tunnel, but this is not available as yet). The Server Side ML-IP device will need to have this network address specified for each tunnel.
- Network B is the IP addressing used on the Internet links (DSL, T1 etc.) connected to the 2344. These can be assigned in the same way normal, non-ML-IP links would be. However, multiple links cannot be on the same IP subnet, unless they are PPP or PPPoE links.
- Network C is the network at the ISP. This must be reachable across the links which make up Network B.
- Network D is made up of tunnel endpoint addresses. These can be drawn from the RFC 1918 address ranges, as no traffic is sent to or from them – they are only ever used by the ePipes for routing purposes. These addresses are assigned at connection time by the server side ML-IP device, and default to using addresses from 192.168.250.1 to 192.168.254.254.

3 Installation and Configuration

3.1 Server Side Unit

We recommend using an ML-IP Concentrator at the server side for all ISP installations. As mentioned above, it should be installed with only one Ethernet interface on the LAN. This requires an installation procedure as detailed below.

We recommend installing the ML-IP Concentrator software with:

- A single NIC configured in the Linux PC, with that NIC configured on a 192.168.n.n subnet. The install script of the software adds routes to any networks attached to the Linux PC at install time, so to avoid conflicts we suggest using a private address. This route will be removed early in the configuration process.
- A single Windows PC with a Java-enabled browser on that subnet for initial configuration (requirements for both PCs are in the Getting Started Guide for the ML-IP Concentrator).
- Another NIC installed (but not configured) under Linux. This will become Ethernet 2 on the ML-IP Concentrator.

Install the software as described in the Getting Started Guide. Once the software is installed and working, configure these items, using the ML-IP Manager (web interface) from the configuration PC:

1. Configure Ethernet 2 with its IP details (in the Network Management->Broadband tab of the ML-IP Manager)
2. Configure Remote Management (Traffic Management->Remote Management) to allow access to the management interfaces from other networks. At install time, the ML-IP Concentrator can only be managed by accessing Ethernet 1 (a virtual address, between the ML-IP Concentrator and the Linux PC).
Because there will not be another PC connected to the unit (only Ethernet 2 will be connected), management on Ethernet 2 must be enabled – this page allows you to specify which IP addresses will be able to manage the box (on any interface). Typically you would add the network which Ethernet 2 is connected to, at least.
Once this change is applied, we recommend you move the configuration PC to that network to ensure management is working, and do the rest of the configuration of the unit from there.
3. Remove the route to the private network used initially for the Configuration PC in Traffic Management->Route Configuration.

You can now start configuring ML-IP Server tunnels in the unit. They will connect to the fixed IP address on Ethernet 2.

There will also need to be some changes made to the routing in the rest of the network, as all packets destined for Network A (the network behind the client side ML-IP device), will need to be routed to the server side ML-IP device. If using NAT on the ML-IP tunnel (not available as yet), the tunnel endpoint addresses (on the client side of each tunnel – these are assigned by the server side ML-IP device) will need to be

3.2 *Client Side Unit*

Configuration of the client side unit is very simple. Once the unit is installed and ready for configuration as per the Getting Started Guide, the Internet Links should be configured (in the Network Connections->Broadband tab) and a tunnel created (in the Multilink-IP->Client tab).

It is important to ensure that the tunnel is set to be the default route of the unit – because ML-IP devices have multiple Internet links of several different types, this can be tricky. If you ensure that during tunnel configuration you tick the ‘Use tunnel as default route’ checkbox, and do not tick the ‘Use this link for incoming traffic’ checkbox on any of the Internet links (this checkbox forces the link to become the default route), the tunnel will become the default route. Do not enter anything in the ‘Configure Static Routes for ML-IP traffic’ step in the tunnel configuration for units which have the tunnel as their default route.

DRAFT

4 Conclusion

ML-IP combines the bandwidth of multiple DSL links between the subscriber and you to increase the speed and reliability of broadband connections, creating a wider pipe into your network at the IP layer. These new multi-megabit services can be further differentiated by speed, range, application and subscriber contention ratios to create a larger portfolio of products.

Using this technology, the optimum service can be made available to every subscriber from small and medium sized businesses, to large enterprises and government. To cater for changing subscriber needs, your ML-IP powered service supports incremental speed upgrades.

This document has been designed to enable ISPs to evaluate or deploy ML-IP services quickly and successfully. Visit our website (www.ml-ip.com) for more information – in particular, the page available from Solutions->ISPs is likely to be of interest to readers of this paper. Please contact ePipe Pty Ltd if you have any questions or comments on this Application Note, or wish to discuss ML-IP solutions.

Web: <http://www.ml-ip.com>

Technical Queries: support@ml-ip.com

General or Sales Queries: info@ml-ip.com

DRAFT

INFORMATION CONTAINED IN THIS DOCUMENT (referred to as an Application Note) IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND BY EPIPE PTY. LTD., EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND/OR FITNESS FOR A PARTICULAR PURPOSE.

The user assumes the entire risk as to the accuracy and the use of this Application Note and the information contained herein. This Application Note may be copied and distributed subject to the following conditions:

- 1) All text must be included without modification and all pages must be included.
- 2) If software is included, all files on the disk(s) must be included without modification.
- 3) All components of this Application Note must be distributed together.
- 4) This Application Note may not be distributed for profit.

Copyright (C) 2003 ePipe. All Rights Reserved.

ePipe is a trademark of ePipe Pty Ltd. All other trademarks are the property of their respective owners.

For further information, contact ePipe by sending email to support@ml-ip.com, quoting the name of this paper in the subject header.

Document Number: AN-EP-013

First Edition: Month, 2002

Keywords: ePipe ML-IP networks ISP

This revision: March, 2003

DRAFT