White Paper

Testing Ethernet Passive Optical Networks

This white paper discusses the different types of Passive Optical Networks with a particular focus on the technology and testing of Ethernet Passive Optical Networks.

Agilent Technologies

Introduction

Passive Optical Networks (PONs) take the "next step" beyond copper DSL and cable to deliver high-bandwidth services to residential and business telecommunication customers. As the name suggests, PONs deliver optical services, or "Fiber to the Home" (FTTH) directly to the end-customer. In some cases, PONs are used to get high bandwidth "close" to the home, enabling higher bandwidth to be delivered over the copper infrastructure over a shorter distance. Regardless of the application, PONs hold the promise of delivering a vast array of high-bandwidth services to the consumer.

Like many technologies, PONs have undergone an evolution in step with the evolving telecommunications industry. Starting with ATM-based PONs (APON), PON technology has evolved to mirror the cost-effectiveness and ubiquitous access of Ethernet: Ethernet PONs (EPONs) are the latest incantation of PON technology, and promise to be widely adopted.

PONs are a shared medium relying upon control mechanisms to manage the allocation of bandwidth amongst the attached end-stations. Being a shared medium, careful testing is required to not only ensure that the infrastructure fairly manages bandwidth across the end-stations, but also to ensure that the underlying mechanisms are suitable to provide Multiplay (video, voice & data) services with a Quality of Experience (QoE) suitable to the provisioned service.

This white paper will discuss the aforementioned topics. First, a discussion of the different types of PONs will set the scene for a discussion of the emergence of Ethernet PONs. This will be followed by a technical discussion of the operation of Ethernet PONs, leading into an examination of the critical test requirements to ensure that an Ethernet PON can indeed deliver Multiplay services.

An Overview of Passive Optical Networking Technology

The Origin of PONs

The price of optical components and fiber drove the development of a shared fiber access system in the 1980s. Utilizing passive optical splitters, multiple end stations could be connected through a shared fiber infrastructure, saving cost in optical transceivers, fiber and installation costs. In spite of being driven by cost requirements, PONs were still too expensive as an access technology and could not compete against cable systems (for video delivery), or copper DSL systems for broadband access.

The plummeting cost of optical components coupled with the competitive requirement for wireline telcos to compete with cable companies for the delivery of video services has breathed new life into PONs. PONs are now viewed as being a cost-effective means for both delivering high-bandwidth services to end-subscribers; or as a means for reducing DSL distances and consequently increasing the bandwidth carrying capacity of DSL services.

How a PON works

A PON consists of a number of Optical Network Units (ONUs), a single Optical Line Terminal (OLT), and one or more passive optical splitters that split the fiber to the ONUs. See figure 1.

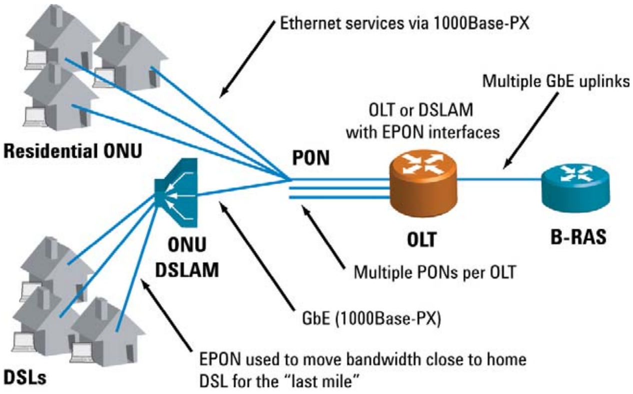

Figure 1: EPON Reference Diagram

As shown in diagram 1, the optical signal can be split multiple times between the OLT and the ONU. The number of times that the signal can be split depends on the optical signal strength and the loss characteristics of each splitter. The ONU may sit in a residence or business (in FTTH applications), or, as also shown in diagram 1, a remote DSLAM is connected to an ONU, and broadband services are distributed via copper. This particular application reduces the distance of a DSL connection, and consequently increases the bandwidth carrying capacity of that DSL connection.

The Passive Optical Splitter

A PON is based on a single component – the passive optical splitter. This component, as its name suggests, is an non-powered optical device that can split an optical signal on an ingress fiber onto multiple egress fibers. A passive optical splitter thus replicates an optical signal across multiple fibers in one direction. In the opposite direction, the passive optical splitter combines signals from multiple sources. For single-wavelength systems, only a single end station can transmit at once, otherwise there will be interference from multiple signals.

Evolving Standards – from APON to EPON and GPON

The first PON architecture was ATM PON (APON), and uses ATM at layer 2. APON rates were defined at 155Mb/s upstream/downstream and an asymmetrical combination of 155Mb/s upstream, and 622Mb/s downstream.

Broadband PON (BPON) was a "rebranding" of APON to remove the stigma of being an ATM only technology. As BPON progressed, it added support for asymmetrical 155Mb/s upstream / 622Mb/s downstream (and, later, 622Mb/s upstream / 1.2Gb/s downstream), and added support for video transmission over a separate wavelength on the PON. Both APON and BPON are defined in the ITU-T's G.983 series of recommendations. BPON has been widely adopted for initial deployments in both Japan and North America.

Gigabit PON (GPON) is another evolutionary step for PON. Defined by the ITU-T in the G.984 series of recommendations, GPON provides backwards compatibility for ATM-based BPON systems while adding new transport support for Ethernet and TDM voice, while adding enhanced traffic rate support, up to 1.25Gb/s downstream / 2.5Gb/s upstream. GPON is seen as the successor to BPON, and is expected to be tapped as the technology of choice for the next round of PON build outs in North America.

Finally, Ethernet PONs (EPONs) promise to offer the ubiquity and price efficiency of Ethernet over a passive optical network. Ethernet PONs are defined in the IEEE 802.3ah Ethernet in the First Mile specification (officially called the "Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks"). Although EPON only transports Ethernet frames, the migration of voice, data and video services over Ethernet (rather than over a dedicated infrastructure) makes a simplified PON architecture based on EPON very appealing. EPON is being widely deployed by NTT and throughout Asia.

Ethernet PONs

Ethernet PONs work in a similar fashion to other PON technologies; that is, the OLT utilizes a protocol to ensure that each ONU transmits only at a specified time, to avoid multiple stations transmitting at the same time and causing interference of the optical signal.

How an Ethernet PON works

The Multi-Point Control Protocol (MPCP) is defined by 802.3ah, and is the protocol that manages access to the PON. A discovery process allows ONUs to register with the OLT. ONUs request bandwidth and the OLT grants transmit windows to each OLT. The fair sharing of bandwidth across multiple ONUs is managed by the Dynamic Bandwidth Algorithm (DBA).

Framing

802.3ah modifies the format of Ethernet frames used in an EPON by redefining the preamble to add fields for a Logical Link Identifier (LLID) plus CRC. The original preamble definition, and the fields modified to support EPON are shown in the following table:

Offset	Field	Standard Preamble/SFD	New Modified preamble/SFD
1	-	0x55	same
2	-	0x55	same
3	-	0x55	0xd5
4	-	0x55	same
5	-	0x55	same
6	LLID[15:8]	0x55	Mode & Upper 7 bits of LLID
7	LLID[7:0]	0x55	Lower 8 bits of LLID
8	CRC8	0xd5	8 bit CRC calculated over offsets 3 through 7

The LLID identifies an ONU on the PON. The LLID is assigned during the discovery process.

Timing Considerations

ONUs may be located far from the OLT – at such a range that the OLT must take propagation delays into account when allocating transmit windows to the ONUs. The ONU and the OLT both maintain 32-bit counters that increment every 16 ns. When the OLT sends a message, it inserts the value of this counter into the message. When the ONU receives the message, it sets its counter to the value received. The ONU places its counter value into messages transmitted to the OLT. When the OLT receives messages from the ONU, it compares the received counter value with its counter value. The difference is the propagation delay between the OLT and ONU. The OLT uses this time difference when allocating transmit opportunities to the ONU.

Discovery

OLTs discover ONUs attached to the PON by sending out a Gate MPCP message (with the Discovery flag set) at regular intervals. The particular interval value is not defined by the standard and is left to the manufacturer. Upon receipt of a GATE message indicating discovery, non-registered ONUs wait a random amount of time before sending Register messages to the OLT to avoid interference between ONUs registering at the same time. After a successful registration, the OLT assigns a Logical Link Identifier (LLID) to the ONU, allowing it to distinguish frames received from, and to target frames at, a particular ONU.

Ethernet OAM

Following registration with the OLT, Ethernet OAM on the ONU starts its discovery process and establishes a connection with the OLT. Ethernet OAM is used on the ONU/OLT link for detecting remote failures, initiating remote loopback, and monitoring the quality of the link. However, Ethernet OAM provides support for "Organization specific" OAM PDUs, information elements and event reporting. Many ONU/OLT manufacturers utilize these organization specific extensions to configure specific capabilities within the ONU. Typically, these extensions configure bandwidth profiles in the ONU to control the amount of bandwidth used by end-subscribers. The non-standard use of Ethernet OAM for these applications is a key test consideration and is a barrier to interoperability between the ONUs and OLTs from different equipment manufacturers.

Authentication and Registration

Also following the registration of the ONU with the OLT, end clients attached to the ONUs authenticate themselves with the network using 802.1x, and then request an IP address with DHCP.

Downstream Traffic

When an OLT has traffic to send to an ONU, it sends the traffic onto the PON, prefixed with the appropriate LLID corresponding to the destination ONU. Because of the point to multipoint nature of a PON, frames transmitted by the OLT onto the PON are broadcast to every ONU. Only the ONU with the correct LLID should forward frames to the attached subscriber. Because downstream frames are sent to every ONU, many NEMs favor Ethernet encryption to secure traffic from hacked ONUs.

One consideration for downstream traffic is video delivery. Due to the broadcast nature of a PON, when a single client subscribes to a video feed, that video feed is broadcast to all stations on the PON. Thus, the amount of downstream bandwidth available to any one client is limited by the amount of bandwidth consumed by all subscribers.

OLTs generally support IGMP snooping, which "snoops" on IGMP join request messages and, when multiple clients subscribe to the same multicast feed, it broadcasts a single copy of packets pertaining to that multicast feed, rather than a copy for each subscriber. This helps limits the amount of traffic on the PON. If an ONU supports multiple subscribes, it may also support snooping.

Upstream Traffic

As discussed earlier, only a single ONU can transmit at a time. The ONU has several prioritized buffers (each corresponding to a distinct QoS level). The ONU requests a transmit opportunity by sending a Report message to the OLT, detailing the status of each of its queues. The OLT will respond with a Gate message, telling the ONU the time of the next transmit opportunity, and the duration of that transmit window.

Management of upstream traffic poses a number of challenges for the OLT. The OLT must be able to manage bandwidth requests from all attached ONUs, and must prioritize transmit grants based both on the priority of the queued traffic as well as balancing requests across multiple ONUs, as dictated by the DBA.

Scalability Considerations for Ethernet PONs

There are a number of scalability considerations for Ethernet PONs, dictated by both the physical scalability of the system as well as considerations based upon the types of services offered over the EPON.

Physical Scalability

Although IEEE 802.3ah does not specify the maximum number of ONUs that can be connected on a single PON, the upper maximum, dependent on the reach to the stations, can be from 16 to possibly as high as 128 ONUs on a single EPON. Each ONU that is added to the EPON requires an MPCP session plus an OAM session. As more stations are added onto an EPON, the risk and impact of a system failure (caused by a power failure, for example) increases significantly. For example, each ONU will need to go through the rediscovery process, register, and then an OAM session will need to be initiated for each ONU – the overall recovery time will vary with the number of ONUs on the PON.

If the EPON is used to shorten the reach of the copper plant, or if an ONU is installed within a multi-tenant unit (MTU), multiple clients can be attached to a single ONU. Consequently, the total number of clients on a single EPON could measure into the thousands, depending on the service provider's physical plant and the services that the provider plans to offer over the EPON.

Service Dimensioning

Service dimensioning refers to the balance that a service provider must strike between service provisioning and physical scalability. For example, even though a single PON could potentially support thousands of clients, a service provider may need to ensure that each client can subscribe to video, voice and high-speed data services. Consequently, physical scalability must be compromised in favor of the high-value services that can be offered over the EPON.

The Path to Interoperable Ethernet PONs

EPON interoperability is a key consideration for service providers planning EPON build-outs. Service providers need the flexibility to provide EPON systems from different vendors (as dictated by price and supply), and will want to ensure that if they create an infrastructure utilizing OLTs from multiple vendors, their subscribers will be able to move their ONUs seamlessly from one location to another.

To emphasize this, the Chinese government has mandated that OLTs and ONUs certified for sale within China must be able to interoperate, ensuring that subscribers can move their ONUs between service providers, ensuring competitiveness amongst service providers.

There are currently barriers to interoperability:
- The Implementation of the dynamic bandwidth algorithm may vary between devices;
- Some implementations use "Organization specific" elements of OAM to set particular behavior (such as traffic management settings) in ONUs;
- Protocol interoperability between implementations of MPCP; and
- Timing issues between different vendors.

Achieving interoperability requires that each of these issues be addressed.

Ethernet PONs and Multiplay Traffic Considerations

An Ethernet PON is an ideal architecture for delivering Multiplay services to consumer customers. The nature of a PON may introduce certain impairments, or may have particular behavior, that can impact service delivery. Three primary factors can affect Multiplay traffic: the nature of the shared medium of EPON, upstream traffic delay, and over utilization of upstream bandwidth.

Shared Medium

As pointed out previously, a PON is a shared medium in both the upstream and downstream directions. Now, for Multiplay services, this can have an impact on services provided over this medium. Take, for example, broadcast video. A subscriber may join a broadcast video stream (via a multicast join). The video server and, eventually, the OLT will stream packets to the subscriber, consuming a percentage of the downstream bandwidth on the PON. If a second subscriber on a second ONU joins a second video stream, it will consume an additional percentage of the downstream bandwidth on the PON. If multiple subscribers join additional video streams, the downstream bandwidth of the PON can be oversubscribed and fully consumed. Enforced policy will dictate how downstream bandwidth on a PON can be shared.

Upstream delay

ONUs must wait for a transmit opportunity from the OLT before they send data onto a PON. Thus, ONUs must buffer upstream traffic, introducing packet latency (delay), jitter and potential packet loss.

Generally speaking, upstream traffic is not as much an issue as downstream traffic. While downstream traffic will consist of revenue-generating services (such as video and voice), upstream traffic will generally consist of low-bandwidth voice and non-revenue generating, but potentially high bandwidth consuming traffic (such as peer-to-peer traffic). Thus for non-priority upstream traffic, a certain amount of loss and latency is generally acceptable.

Upstream packet latency is variable and depends both upon the number of other ONUs on the PON, and the amount of data that they have to send. If all ONUs have the same amount of traffic to send, and are thus allocated sequential transmit opportunities, packet latency can be calculated as a function of the allocated transmit duration times the number of ONUs with traffic to send. However, considering the dynamic bandwidth algorithm manages the allocation of bandwidth, the actual packet delay may be greater or less than this nominal value. Excessive upstream latency can cause subsequent packet loss. Excessive packet loss can in turn have a negative effect on subscriber services, such as voice or significantly reduced throughput for file sharing applications, negatively influencing a subscriber’s opinion of the offered service.

Over-Utilization of Upstream Bandwidth

Upstream traffic quality is greatly affected by the amount of upstream traffic sent by other subscribers. Effective management of the traffic profiles of each subscriber will be paramount to ensure an equally accessible PON. For example, a subscriber involved in file sharing may be uploading a large, constant stream of data (without his explicit knowledge, as these transfers take place in the background). Considering that file sharing constitutes (by some reports) up to 60% of consumed bandwidth on the Internet, this type of traffic needs to be carefully managed.

Testing Ethernet PONs

This section details some key test criteria for testing Ethernet PONs. This is not a comprehensive "test plan" for testing an Ethernet PON, but rather highlights some specific test criteria to test the most critical components of EPON.

An Overview of EPON Testing

Standard Test Configuration

For these tests, it is necessary to generate traffic through the Ethernet PON and to have a means of observing the packet flow within the PON. The following diagram shows the test configuration:

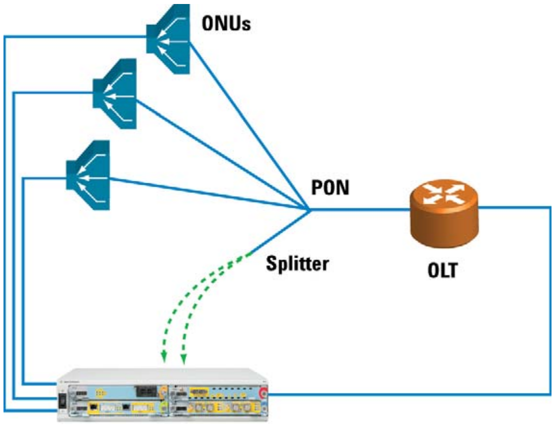

Figure 2: EPON Test Configuration

Traffic generators connected to the Ethernet PON generate client or network traffic. These traffic generators must encapsulate a test payload (containing a time stamp and sequence number) into each packet. With this test payload, packet loss and latency can be measured at either end of the EPON.

The EPON monitor, located between the passive optical splitter and the OLT "snoops" packets flowing between the ONU and the OLT. It must be capable of processing data packets on the EPON, extracting the test payload, and displaying packet loss and latency statistics. The EPON monitor must also have trigger, filter and capture capabilities to trigger on MPCP, OAM or data packet events and capture and decode the corresponding protocol transactions.

Key Tests

MPCP Registration

Objective:
In this test, the objective is to ensure that an ONU registers with an OLT within the correct time window.

Discussion:
It is important to ensure that an ONU is registering within the specific time window allocated to the discovery process.

Methodology:
To perform this test, it will be necessary to ensure that:

- The OLT sends out a Gate MPCP message with the discovery flag set at regular intervals;
- The ONU attempts registration only during the discovery interval; and
- The ONU attempts registration after pausing a random amount of time before sending a Register message.

The test configuration outlined at the beginning of this section is used. One ONU is the target device to test for registration. The EPON monitor is configured to trigger on the Gate message indicating discovery from the OLT, and to capture the subsequent exchange when the ONU sends its Registration message to the OLT.

The key measurement is to measure the time at which the ONU registers with the OLT. Having captured the protocol exchange, the Gate message can be decoded and the window determined. This window can then be compared to the timestamp of the received Register message to ensure alignment within the discovery window.

This measurement is repeated to ensure that a random interval is used by the ONU from the start of the discovery window through to when it transmits the Register message. Repeated measurements will ensure that a random amount of time passes between the opening of the discovery window and the Register message.

Downstream Traffic Broadcast

Objective:
In this test, downstream traffic is addressed to specific ONUs. Although, by the nature of a PON, the traffic is "broadcast" to each ONU, only traffic specifically addressed to that ONU should be forwarded to the end client.

Discussion:
After the ONU has registered with the OLT, downstream traffic is offered to the network side of the OLT, and the traffic is forwarded to the appropriate ONU, prefixed with the appropriate LLID.
The key measurement is to ensure that downstream traffic is affixed with the appropriate LLID on the EPON, and to ensure that only the target ONU forwards traffic to its client.

Methodology:
Traffic is offered to the network side of the OLT with a destination address corresponding to a client behind a target ONU.

Traffic should be forwarded onto the PON, prefixed with the appropriate LLID. This is confirmed by capturing data packets on the PON, and ensuring that the LLID is correct. At the client side of the ONUs, traffic should be forwarded from the target ONU, and not from any other ONU.

Upstream Transmit Alignment

Objective:
This test verifies the alignment between the transmit window granted by the OLT to the packets transmitted from the ONU.

Discussion:
This test requires a high degree of precision to accurately measure the time of packets transmitted by the ONU, and to ensure that these packets are transmitted within the specific time window granted to the ONU by the OLT through a Gate MPCP message.

To measure this, it will be necessary to trigger upstream data capture (from the ONU to the OLT) on the Gate message sent by the OLT to the ONU. The Gate message will be captured, as well as the data upstream data and Report messages from the ONU.

Methodology:
On the downstream monitor port, set a trigger to start bidirectional capture when a Gate message is sent to the LLID corresponding to the target ONU. Also, set a capture filter so that only upstream packets from the target ONU are captured.

Send data packets into the target ONU. A Report message will be sent from the ONU to the OLT requesting a transmit window. The OLT will respond with a Gate message. This message will trigger capture to start, and will capture both the Gate message, as well as the upstream data packets from the target ONU to the OLT.

After capturing data, the timestamps in captured Report messages (from the target ONU) is compared to the timestamp in the Gate message (from the OLT to the ONU). The times can be correlated to determine whether the ONU sent it's data during the specific window.

Upstream Traffic Prioritization

Objective:
This test ensures that priority packets are delivered ahead of low priority packets from the ONU to the OLT.

Discussion:
ONUs maintain several queues for different traffic priorities. For example, voice packets should be prioritized ahead of data packets (which may contain lower priority web or peer-to-peer traffic). Packet priorities are distinguished by the VLAN priority field. An EPON can also be configured so that the packets corresponding to different VLAN priorities are mapped directly onto a distinct LLID.

The ONU can be forced to prioritize traffic by sending data into it that exceeds the bandwidth available on the EPON. This can be accomplished by transmitting priority traffic from a second ONU on the EPON at a high bandwidth, restricting the bandwidth available for the target ONU, and forcing the ONU to prioritize traffic.

In essence, this is a test of one aspect of the dynamic bandwidth algorithm – ensuring fair access to the EPON amongst a number of ONUs, and ensuring that high priority traffic receives the appropriate resources for transmission.

Methodology:
Low and high priority packets are sent into the target ONU. High priority packets are sent into a second ONU at a high bandwidth.

The packet loss and latency for high and low priority packets from each ONU can be measured at the OLT. There should be minimal impairments on the high priority stream – any impairments should appear in the low priority streams.

It is useful to observe the actual location of packet loss and latency. Using the monitor ports, packet loss and latency can be observed directly on the EPON. Packet loss & latency for low priority traffic streams should be occurring directly at the target ONU.

Upstream Traffic Sharing

Objective:
This test ensures that one ONU cannot consume all of the traffic on an EPON.

Discussion:
The dynamic bandwidth algorithm is meant to ensure that all ONUs have fair access to the EPON – that is, one ONU (or, several ONUs) should not be able to consume all of the bandwidth and “starve” other ONUs.
This test generates the conditions under which a client attached to one ONU attempts to consume all available bandwidth. A second ONU, attempting to transmit a nominal amount of packets should be able to transmit those packets.

Methodology:
A high volume of packets is sent into one ONU (or a group of ONUs), such that all bandwidth on the EPON is consumed by this traffic.

On the target ONU, packets are sent into it at a nominal rate – within a bandwidth profile assigned to that ONU.

Packets received at the OLT from the target ONU should experience packet latency and minimal packet loss within the boundaries acceptable for a fully subscribed PON. Packets received at the OLT from the other ONUs should experience a higher packet loss and latency.

Single Copy Broadcast

Objective:
This test verifies that an OLT implement "single copy broadcast" for sending multicast traffic to multiple subscribers.

Discussion:
Single Copy Broadcast (SCB) is used to reduce bandwidth consumption on an EPON when multiple subscribers join a single multicast feed. For example, several clients may join the same broadcast channel. Instead of the OLT sending individual packets to each ONU for each client that has joined that multicast group, a single copy can be sent (indicated by setting the Mode bit in the 802.3ah preamble to "1"). Receiving ONUs will forward the packet based on the MAC destination address.

It is difficult to verify that SCB is working correctly without observing the actual packet flow on the EPON. If an OLT is not implementing SCB, there is no way to verify this from observation of the end-to-end traffic – a monitor port can ensure that only a single copy of a multicast packet is sent, and that the Mode bit is set correctly.

Methodology:
One ONU subscribes to a multicast feed (by sending an IGMP Join message).

The tester connected to the OLT should send multicast traffic to the ONU.

A second ONU joins the multicast feed.

The OLT should "snoop" the second IGMP Join message, and should send a single copy of all multicast traffic with the Mode bit set to "1". The downstream monitor port is set to trigger and capture downstream packets with the Mode bit set to "1". The captured packets are inspected to ensure that there are no duplications of packets – that is, only a single copy is being sent to all destinations.

Summary

In summary, Ethernet PONs offer a simplified mechanism for supporting Multiplay services by extending optical access to businesses and consumers. By nature, devices connected to a PON have to share upstream and downstream bandwidth – the dynamic bandwidth algorithm (DBA) uses the Multi-Point Control Protocol (MPCP) to grant transmit windows to the ONUs.

An EPON must be carefully tested to ensure that it can support the services required by service providers, and to appropriately dimension EPON access networks to support those services. These tests cover both downstream utilization as well as ensuring that upstream bandwidth is mapped correctly to individual ONUs, and shared fairly across multiple ONUs.

www.agilent.com/find/n2x

Sales, Service and Support

United States:
Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:
Agilent Technologies Canada Inc.
2660 Matheson Blvd. E
Mississauga, Ontario
L4W 5M2
1-877-894-4414

Europe:
Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-2323

United Kingdom
07004 666666

Japan:
Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:
Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 269-7500
Fax: (305) 267-4286

Asia Pacific:
Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 3197-7777
Fax: (852) 2506-9233

Australia/New Zealand:
Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

Printed on recycled paper

Printed in USA October 01, 2008
5989-4070EN

Agilent Technologies