The Ethernet Effect: Collaboration, Interoperability and Adoption of New Technologies

A University of New Hampshire InterOperability Laboratory White Paper

In Collaboration with Dell'Oro Group

Including a Case Study of Japanese Service Provider
Nippon Telegraph & Telephone (NTT) Group

April 2006

Executive Summary

Ethernet's adoption by global data-communications service providers, enterprises and the home networking market has spurred the evolution of the simplest, most basic building block of networking into a far more complex technology. Beneficial new features, including increased speed, capacity, and power, have provided a means for Ethernet migration and deployment in new markets. This migration, however, has brought into question Ethernet's ability to provide a solid user experience and properly deliver quality services in these new configurations. Detailed interoperability and conformance testing is underway to reassure early adopters and provide feedback to the standards bodies and equipment vendors in an effort to facilitate the technology adoption cycle.

The University of New Hampshire InterOperability Laboratory (UNH-IOL) has been evaluating Ethernet equipment since 1990. This white paper represents the first public summary of objective, multi-vendor Ethernet conformance and interoperability testing conducted over an extended period. This whitepaper summarizes seven years of Fast and Gigabit Ethernet testing data gathered through the end of 2005, including test reports from 900 individual products representing several hundred companies. NTT has contributed a carrier's perspective of the benefits this information provides to commercial network deployments. Dell'Oro adds context to this information with statistics on Ethernet port shipments.

Introduction

Since Robert Metcalfe's invention in 1973, Ethernet has remained a rare constant in a rapidly transforming communications industry of disruptive innovation. As can be seen from data provided the Dell'Oro Group, not only have Ethernet switch port volumes risen dramatically, but each successive iteration of the technology has or is projected to surpass the prior speed.

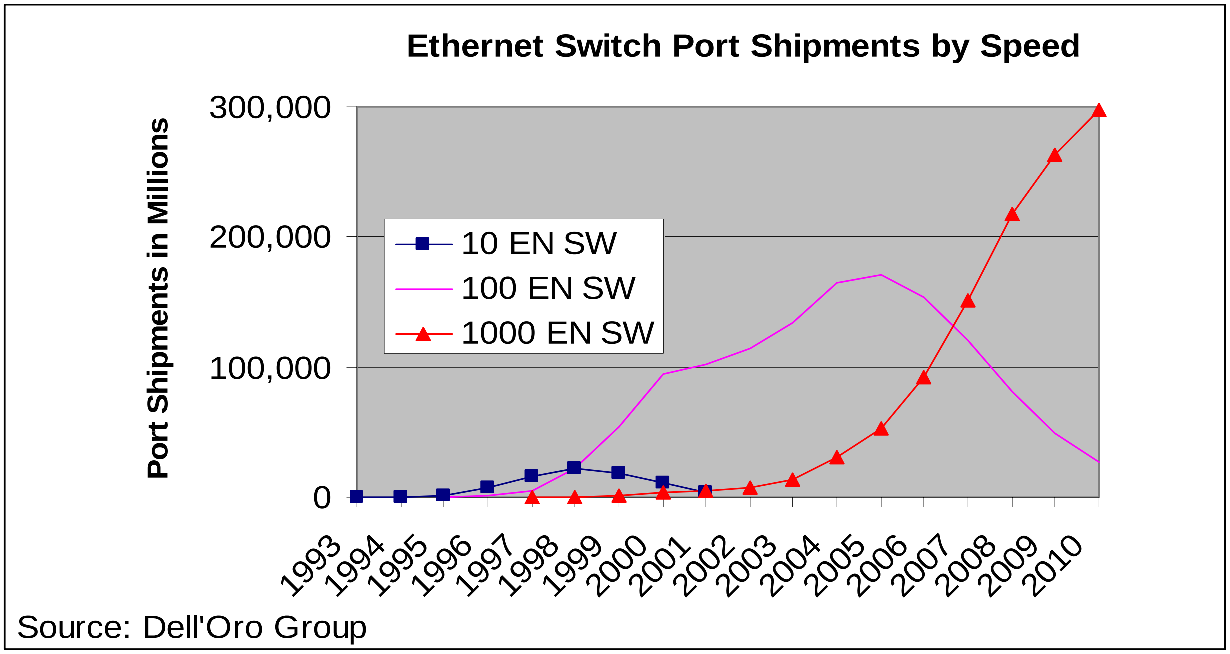

The original patent describes Ethernet as a "multipoint data communication system with collision detection." While it has remained, as originally described, a system allowing, "distributed packet switching for local computer networks," Ethernet has long outgrown its role as a system for connecting computers within a building using hardware running from machine to machine.

Ever since Metcalfe convinced Digital Equipment, Intel, and Xerox Corporation to collaboratively promote Ethernet as a standard, industry cooperation and laboratory testing have fostered Ethernet's success in ever more complex multi-vendor and multi-platform environments. Today, Ethernet is the indispensable building block of global networking and the Internet. If Ethernet and the Internet protocol comprise the neural framework, interoperability allows the neurons to communicate.

Ethernet is already the most dominant transport technology in the enterprise, connecting servers, workstations and printers. Network operators are beginning to deploy Ethernet in their core and access networks. With the maturation of IP-based applications, Ethernet is increasingly used in core telecommunication networks as a supplement or replacement for older carrier facilities, such as asynchronous transfer mode (ATM) and other legacy transport technologies. In access networks, Ethernet is replacing traditional phone and digital subscriber lines, providing high-speed connectivity to homes and businesses. This new network infrastructure has the poetential to enable high-bandwidth applications such as the "triple play" of voice, video and data for residential services.

Since 1990, the UNH-IOL has evaluated Ethernet interfaces for hundreds of companies. The UNH-IOL works closely with the standards committees, systems vendors and component suppliers. Technologies include 10Base-T, Fast Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet, and Ethernet in the First Mile (EFM). Overall UNH-IOL generates roughly 200 Ethernet interoperability reports each year. This white paper, the first public study of this scope, draws upon data gathered from seven years of continuous Ethernet testing. It summarizes and analyzes the findings from approximately 900 individual product test reports issued to several hundred companies.

Technical Background

Interoperability is often a broad and intangible concept. The components of interoperability may include multiple aspects, including system sub-components, physical/logical interfaces, software-based protocols and network-level applications. When performing an investigation of interoperability, a very clear and precise definition must be used to clarify the scope of the project. For the purpose of this paper, interoperability is defined as the completion of Ethernet Auto-Negotiation and the establishment of a functional and operating Ethernet link between two link partners.

The following three cases illustrate different initialization scenarios that may exist within a network under test:

Case 1 – Both link partners are powered off and then connected with a compliant channel (fiber or copper). One of the link partners is powered on and is allowed to fully initialize before the second link partner is powered on.

Case 2 – The same as Case 1, but the order in which the link partners are powered is reversed.

Case 3 – Both link partners are powered off, then powered on and allowed to initialize before the compliant channel is connected between them.

In all three instances, both link partners are required to complete Auto-Negotiation in order to establish the link. The test procedure forces the link partners to transmit a specified amount of network traffic between them to verify the link status. If the devices establish a link and can pass traffic between them, then the particular case is said to "pass." If the devices are unable to establish a link or can establish a link but not transport traffic, then the particular case is said to "fail." For two link partners to "pass" the entire interoperability test, all three test cases must be passed. A failure in any one case signifies an overall failure of interoperability. It is important to note that even though two devices may pass the standard of interoperability set forth in this paper, it does not guarantee that the two devices will pass all definitions of interoperability or that the two devices will always be interoperable under all conditions.

This paper will examine two compliance tests in addition to interoperability. The first deals with reception of preamble, the second with the value of an Auto-Negotiation timer. These two tests were selected because of their significance for service providers with the understanding that they cannot be used to illustrate full compliance (or non-compliance) of any of the devices tested.

The reception of preamble test examines whether or not a device will accept frames that contain fewer or extra bytes of preamble than would normally be expected. Depending on the technology or circumstances, it is possible for the preamble on a frame sent by an originating station to be modified by the time it finally makes it to the destination. Generally speaking, devices should be insensitive to the reception of preamble, as it is simply discarded at the receiver. This is a test performed during Media Access Control (MAC) compliance testing.

The link_timer test, a 1000BASE-X Auto-Negotiation compliance test, determines the value that the device under test (DUT) uses. The timer has a nominal value of 10ms, with a range of +10ms and –0s. This timer is used several times during the Auto-Negotiation process, and had a significant effect on the length of time the Auto-Negotiation process takes to complete.

This paper primarily surveys and analyzes specific results from 1999 – 2005 as obtained by engineers in the UNH-IOL Gigabit Ethernet and Fast Ethernet Consortiums. Specific companies and products are not listed in the results. This is to maintain the individual companies' confidentiality agreements.

1000BASE-X Interoperability Testing

The IEEE published the standard defining 1000BASE-X, IEEE 802.3z in June of 1998. Shortly after, the first large-scale multi-vendor interoperability testing event was staged at the UNH-IOL. Consortium based testing began near the end of 1998 and grew rapidly through 1999. The figure shown below shows the results of interoperability testing done by the Gigabit Ethernet Consortium (GEC) at the UNH-IOL from 1999 – 2005.

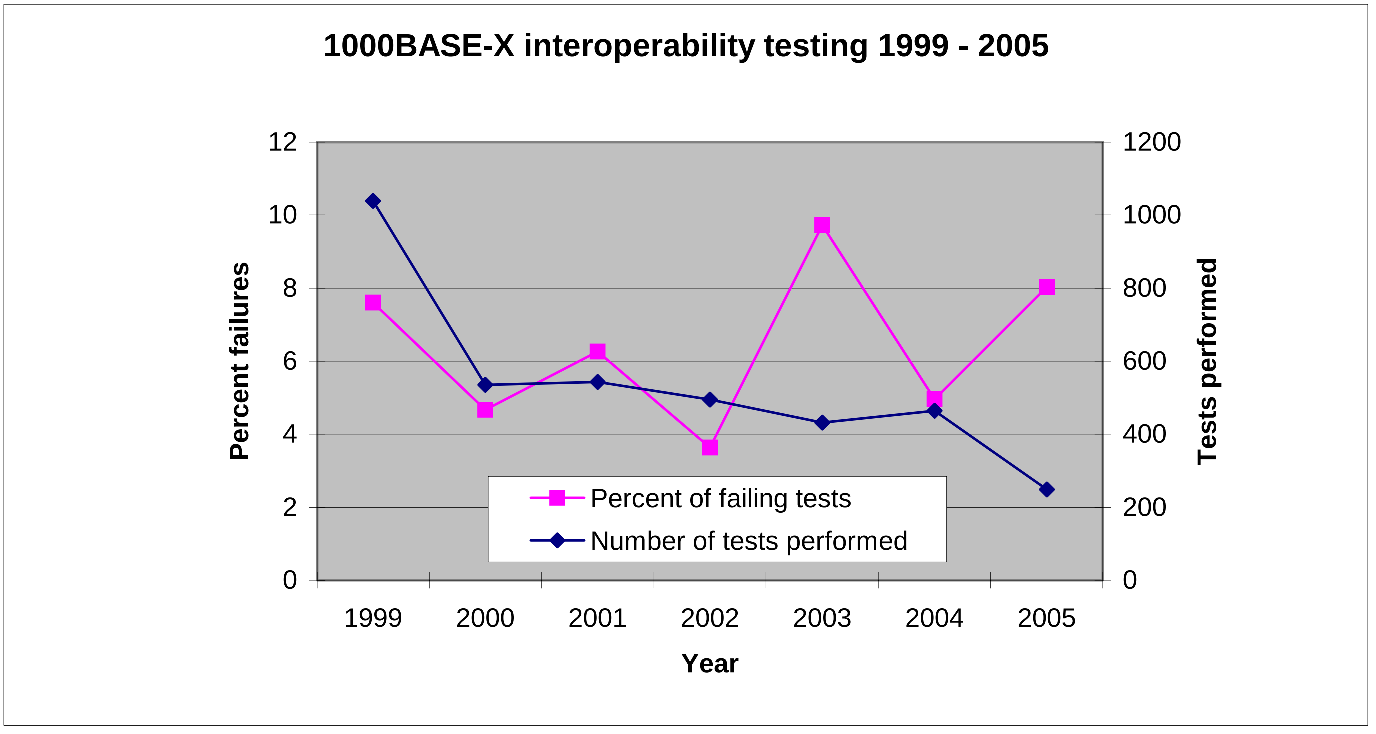

There are a number of interesting details in this figure. First, it is clear that industry has demanded a fairly constant level of 1000BASE-X interoperability testing. The reason that 1999 appears to have done significantly more testing is that during the first year of GEC operation, interoperability testing was performed on all available link partners, each of which was required to loan one piece of equipment to a shared interoperability test bed. As the test bed grew in size, a new testing process was put into place to handle the scale of testing. Beginning in the year 2000, the number of link partners was limited to 20 devices.

The other primary component of interest is the failure rate of interoperability. On average, the failure rate has been approximately 6%. There appears to be no definite trend leading towards complete interoperability, suggesting that a small amount of failures will continue to be present in future years. However, 1000BASE-X has shown itself to be a very successful technology, widely deployed all over the world in a variety of networks and supporting multiple applications.

For such a successful technology, the question must arise as to why there appears to be an interoperability floor, and why the technology cannot become completely interoperable. In order to fully understand the reasons, additional detail must be given regarding the type of testing that is performed. Typically, a vendor will provide a product to the GEC and will select 20 devices to test against for interoperability. It is not uncommon for a device that is suffering from compliance problems to also suffer from interoperability problems, as there is a relationship between the two. For example, in 2005, one of the prototype devices that underwent interoperability testing had numerous compliance problems and ended up failing interoperability testing with all 20 link partners that it was tested against. Often times, once the network component or system vendor is made aware of these results, the company's engineers will attempt to fix the problems and return for additional testing. Thus interoperability and conformance issues will continue to exist as long as new prototype devices are developed. Many of these compliance and interoperability problems, however, are resolved before the device is made commercially available.

1000BASE-T Interoperability Testing

The standard defining 1000BASE-T, IEEE 802.3ab, was published in June of 1999. The UNH-IOL began to see early implementations of 1000BASE-T products early in 2000. Initially, only one silicon solution was available, and therefore the first set of products all contained the same chips. About a year later, silicon from multiple sources started to hit the market, and currently, a large number of both silicon and system vendors have released 1000BASE-T products. The figure shown below depicts the results of 1000BASE-T interoperability testing done from 2000 – 2005.

Of immediate interest is the high amount of 1000BASE-T testing that has been requested and the rate at which it is increasing. Clearly, the industry has placed a great importance on 1000BASE-T interoperability testing, and there are no signs of the demand slowing down. As mentioned with 1000BASE-X, each device in for testing is allowed 20 different link partners. Therefore, it is clear that the number of devices arriving for testing is also increasing.

Another important trend in the 1000BASE-T data is that the results appear to be approaching a level that would allow only a very small percentage of failures. Apart from the first year, when only one vendor's silicon was available, and the second year, when multiple implementations first began to arrive, the interoperability results have taken a very significant and positive turn. Over each of the last five years, there has been a marked decrease in the number of interoperability failures, compounded by the fact that the number of interoperability tests has continued to increase. With less than one percent failures observed between 2004 and 2005 over the course of almost 2300 tests, it seems clear that 1000BASE-T has an almost perfect level of interoperability.

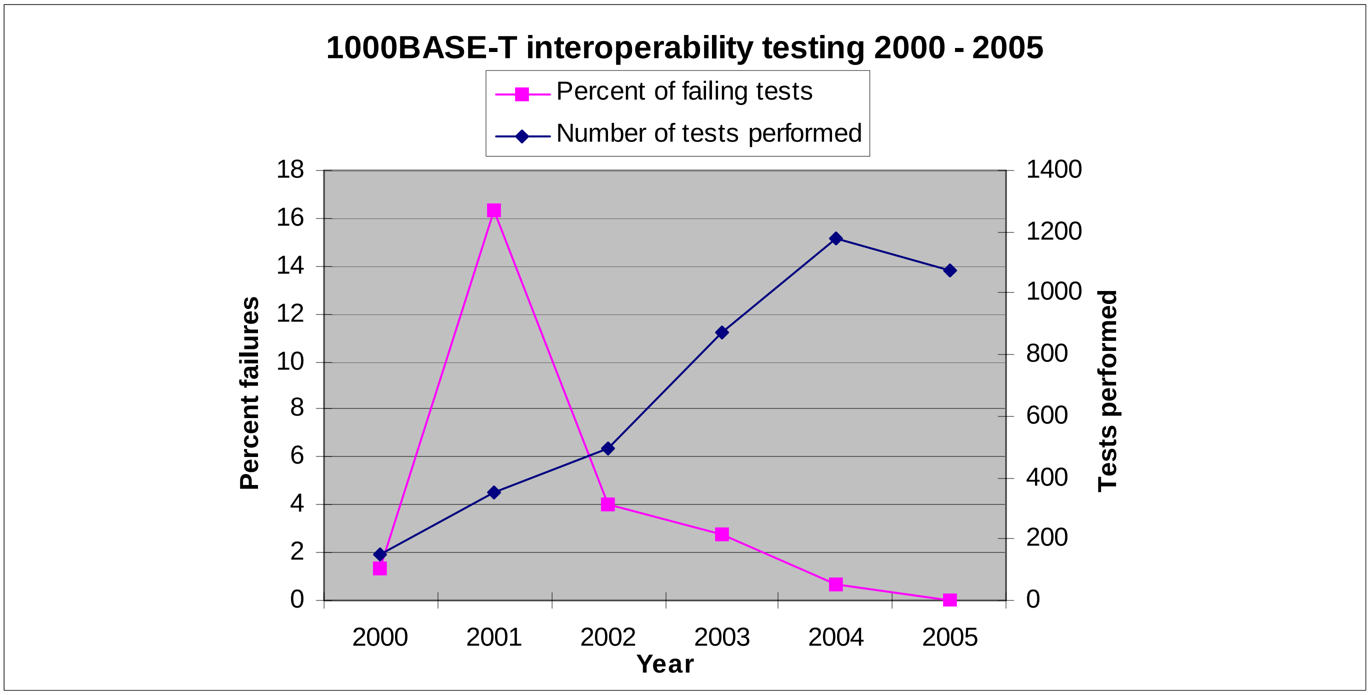

The question must now be asked as to why the results for 1000BASE-T are so different than the results for 1000BASE-X. Part of the difference can be explained as a function of the amount of testing performed. In 2005, the GEC performed just over 200 1000BASE-X interoperability tests, which is approximately five times less than the number of 1000BASE-T tests. Therefore, one device having compliance and interoperability problems with 20 link partners for 1000BASE-X would represent a failure rate of 10%, whereas the same number of failures would represent a failure rate of just 2% for 1000BASE-T.

Reception of Preamble

The results from the reception of preamble test are shown in the figure below for both 100BASE-TX devices and 1000Mb/s devices (both 1000BASE-T and 1000BASE-X). It is clear, in all technologies, that the current trend shows devices becoming more and more insensitive to the reception of preamble of arbitrary lengths.

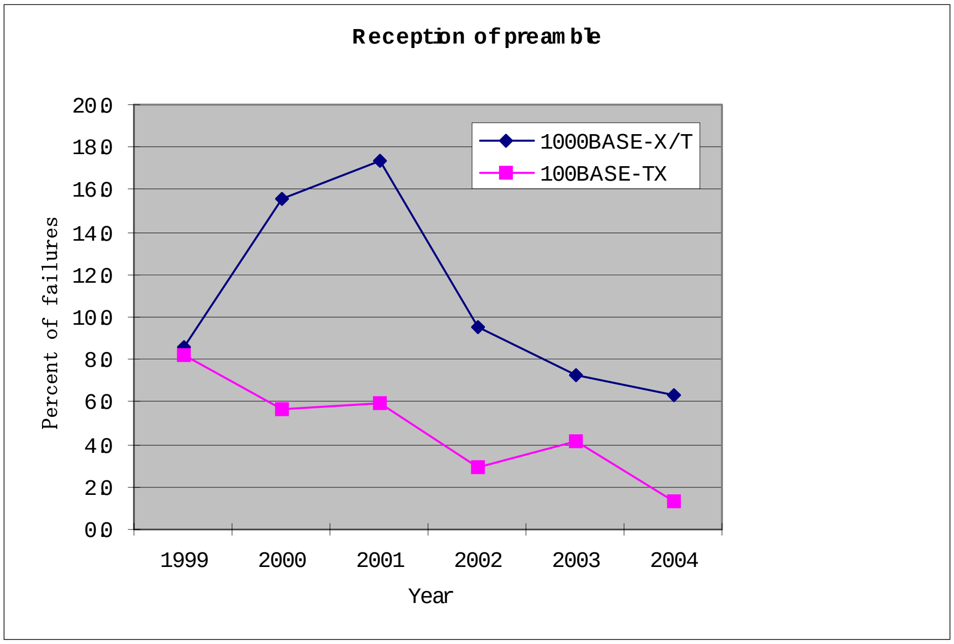

Distribution of link_timer

The following two graphs depict the distribution of the Auto-Negotiation link_timer value for 1000BASE-X devices. The graph on the left shows the total distribution, and the graph on the right is zoomed in slightly, ignoring the 200ms and 50ms values from 2001 and 2002, respectively. It is clear that over time, the value of link_timer has converged to two main camps, one around 10ms, and one around 15ms, with a higher concentration at the lower end. This de-facto standard that industry has arrived at will reduce the amount of time needed to re-establish a broken link.

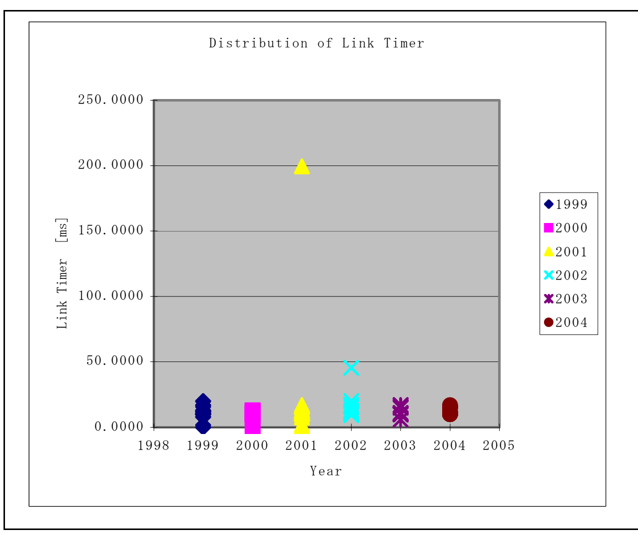

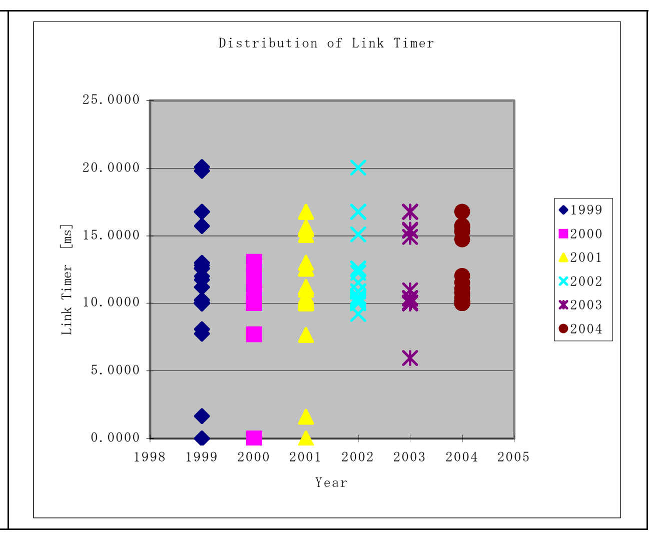

Case in point: The Service Provider Perspective

It has become increasingly common for communication carriers to construct their networks using a variety of commercial equipment from an ever wider variety of suppliers. Vendors manufacture this equipment based on standard documents and proprietary technologies but do not always take into account all of the carriers' specifications. As a case in point, Nippon Telegraph and Telephone (NTT) group, Japan's largest telecommunication carrier, is currently extending metropolitan and residential network services using Ethernet technology. NTT regards Ethernet as one of the most promising mediums in building out network infrastructure, because the technology has come to realize high-speed communications and long-distance transmissions as well as its conventional feature: cost-competitiveness.

NTT Service Integration Laboratories (NTT SI Lab), an R&D division of NTT, routinely conducts extensive testing and verification of network services prior to their commercial release. This process ensures NTT can deliver carrier-class quality services to its customers. During the technological verification processes, NTT SI Lab frequently finds a certain number of issues. In analyzing these issues, NTT SI Lab has noted that a significant portion of these problems can be construed as a lack of interoperability, often as a result of ambiguity in the standards documents in the Ethernet layer (OSI Layers 1 or 2). Many additional issues may be found at an upper layer level (OSI Layer 3 or higher). The problems that specifically derive from the Ethernet layer tend to arise during the final stages of the verification processes for commercial release of a network service. These issues are often difficult or time consuming to modify and thus have a negative impact on the verification schedule.

As a result, NTT SI Lab has become more attentive in Ethernet specifications, increasing awareness of issues that may cause interoperability problems. Auto-Negotiation interoperability, checking preamble length and the distribution of link timer have certainly been responsible for some existent troubles noted in the NTT SI Lab. The IEEE specifies possible ranges of values for some parameters of the Ethernet protocol, such as link timer value. The choice of the value, however, is up to the vendor implementing the standard. Conformance therefore does not ensure interoperability; differing interpretations may cause unexpected issues or degradation in service quality when devices from multiple vendors are interconnected. This becomes a matter of serious consideration for carriers such as NTT. Much of the UNH-IOL's efforts in this area are focused on detecting ambiguous points in the Ethernet specifications as revealed in interoperability testing, and to accelerate interoperability among various vendors. The organization seeks to do this both by recommending changes to the standard and by furnishing reports to individual vendors pinpointing issues and points of failure.

In establishing a carrier-class service using multi-vendor equipment, it is desirable to determine these kind of parameters as specifically as possible in order to provide stable services. One way to do this is to rely on the so-called "de facto standardization" process. Each equipment vendor bringing Ethernet devices to test through the UNH-IOL receives confidential, detailed results. The mere claim of interoperability does little or nothing to suggest which specific values within which specified range enable interoperability before carriers connect such devices. However, it is hoped that exhibiting the facts in reports and white papers such as this one can further accelerate the progress of Ethernet interoperability in the field.

Conclusion

This paper presents data collected over seven years of gigabit Ethernet interoperability testing performed at the University of New Hampshire InterOperability Laboratory. It is clear from the data and market information that the successes of gigabit Ethernet, along with the interoperability of the technology go hand in hand. As can be seen from the below figure, both 1000BASE-X and 1000BASE-T are highly successful technologies, with Dell'Oro estimating 35 million and 80 million ports deployed to date respectively worldwide.

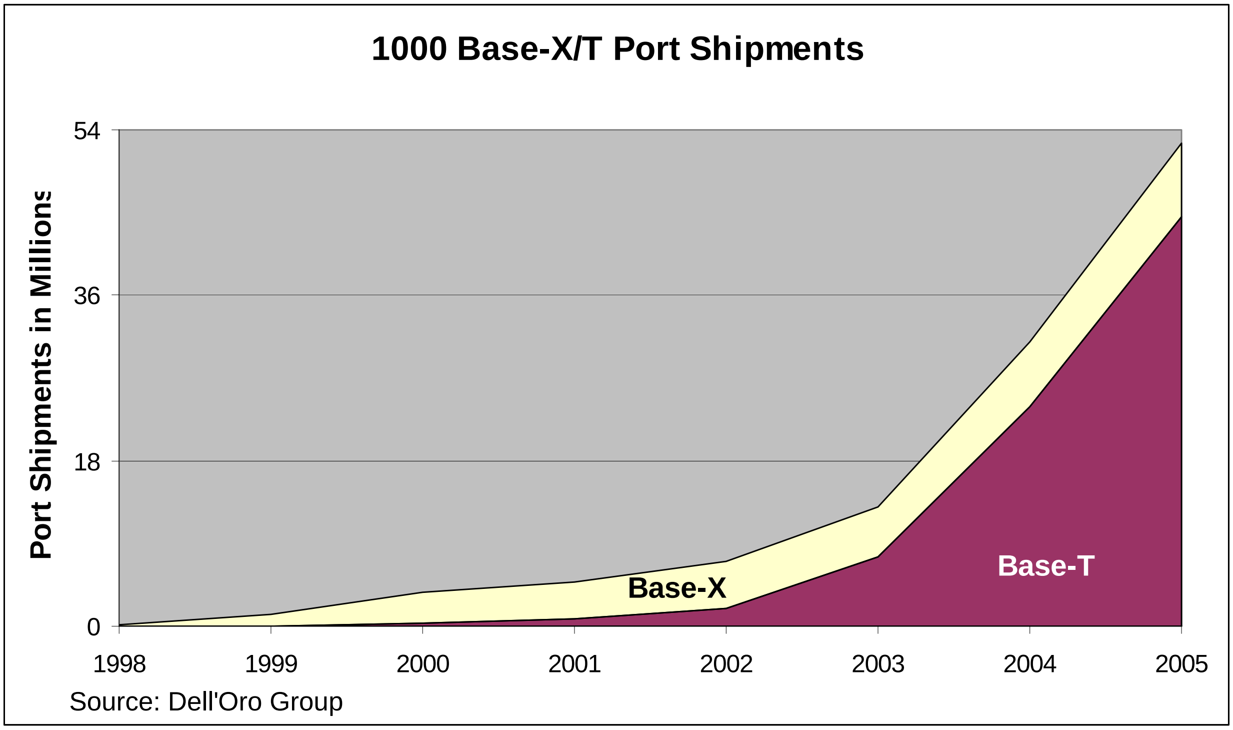

The success of these technologies, and the rate of continued deployment by service providers, must be at least partially attributed to comprehensive and thorough compliance and interoperability testing.

Special thanks to:
Seamus Crehan, The Dell'Oro Group
George Kimura, NTT Service Integration Laboratories
Eric Lynskey, UNH InterOperability Lab
Ben Schultz, UNH InterOperability Lab
Chris Volpe, UNH InterOperability Lab