napatech

NEXT COMPUTING

February 2011

Traffic Generation for the Mainstream Ethernet Market

Testing can't always be performed remotely and issues can't always be seen at a distance. Hands-on testing in the field is often vital for finding a solution quickly. But, to do this effectively, you need the right tools with enough power and performance to deal with whatever challenges you might meet. The combination of NextComputing's portable PC servers and Napatech's intelligent network adapters provides a powerful, yet practical tool for field testing. NextComputing's servers have the power and performance to support the testing applications you need, while Napatech's intelligent network adapters provide both full line-rate transmission and reception up to 20 Gbps. Ideal for network test and forensic activities in the field.

Traffic Generation for the Mainstream Ethernet Market

WHITE PAPER

A NEW VISION FOR TEST AND MEASUREMENT

This white paper provides a new vision for test and measurement equipment vendors: Make sure that every engineer who needs a test and measurement system can get one! The availability of commercial off-the-shelf hardware and software now makes it possible to build affordable test and measurement systems that can help test and measurement equipment vendors to broaden their market to include development engineers who are normally deprived test resources due to cost. To take advantage of this opportunity, test and measurement equipment vendors need to re-think their product development strategy for the mainstream market.

TEST AND MEASUREMENT: A HIGH-RISK BUSINESS

High-tech development is inherently a risky business. Significant investments need to be made up front on the promise of customer adoption and volume sales in the future. High-tech companies constantly weigh the advantages of being the first mover in a market versus the risks of backing the wrong technology horse. For this reason, most high-tech companies adopt a cautious approach to new technologies.

Test and measurement equipment vendors do not have this luxury. When other high-tech companies make their technology bets, they expect test and measurement equipment vendors to already have the test equipment available to support their development. In other words, test and measurement equipment vendors need to make their bets first!

For example, consider 40 Gbps and 100 Gbps Ethernet. Test and measurement equipment vendors had to be the very first companies to implement these technologies, ahead of available standards in order to provide the basis for other technology companies to evaluate the technologies. Test and measurement equipment vendors simply do not have the luxury of a "wait and see" approach.

HIGH-RISK DEVELOPMENT

Test and measurement equipment vendors are therefore inherently pursuing a high-risk development model:

- The market for the new technology is unclear
- Standards are often not available
- Components are often not available
- In-house development is therefore a requirement for many parts of the solution.

To add to this risk, sales of test and measurement equipment will be affected by the price level. What is the right price level to ensure market adoption and adequate volume, while at the same time ensuring cost coverage?

It is therefore no wonder that test and measurement equipment is often considered expensive by end users, but with consideration of the above, it is clear that it cannot be any different.

Or can it?

THE EARLY ADOPTER AND MAINSTREAM MARKET

It is certainly true that when pursuing new, uncharted technology territory that the test and measurement equipment vendors are pursuing a high-risk business model that needs coverage. But, what happens when a technology becomes mainstream? What is the right business model to pursue in such circumstances?

From Geoffrey Moore's seminal work on "crossing the chasm", we know that there is a difference between addressing the needs of the early adopter market and crossing the chasm to meet the needs of the larger mainstream market. Moore's model goes so far as to suggest that a different business model and different products are required.

The reason for this is that early adopters are usually on the bleeding edge of technology development themselves and more interested in achieving a performance boost or other competitive advantage over other technology vendors. They are therefore willing to pay a premium for new technology solutions.

The mainstream market, on the other hand, is more practical and focused on value for money. They are not interested in a performance or other competitive advantage at any cost, but at the right cost and at low risk.

Test and measurement equipment vendors have addressed these concerns using standardization, where these companies are often the most active. However, this is often counterproductive, as it actually commoditizes their products (i.e. if all products are built on the same standards, the only distinction is price). If price becomes the main competitive parameter, how can test and measurement equipment vendors ensure that they can win in the mainstream market?

The answer is to re-think how test and measurement equipment is designed, developed and manufactured for the mainstream market.

RE-THINKING TEST AND MEASUREMENT EQUIPMENT DEVELOPMENT

The traditional design, development and manufacturing model for test and measurement equipment is in-house development, and naturally so. It is necessary if one needs to stay at the absolute forefront of technology developments.

The flaw in this approach is in its application to not only bleeding-edge technology, but also mainstream market technologies. As mentioned earlier, customers in the mainstream market are more cost-conscious and will either buy less than they need and/or shop around for the cheapest offer. The latter is made easier by standardization.

Maintaining an in-house development and manufacturing model for the mainstream market is costly, but also unnecessary. The factors that motivate for an in-house development model in the early adopter market are no longer relevant in a mainstream market:

	EARLY ADOPTER MARKET	MAINSTREAM MARKET
Market potential	The market for the new technology is unclear and needs are not well understood.	The market for the new technology is measurable and needs are well understood.
Standards availability	Standards are often not available.	Standards are ratified.
Product availability	Few products are available and components are often not available.	There are many vendors of products at various levels of the value chain.
Premium justification	Customers are willing to pay a premium for performance or other competitive advantage.	Customers are focused on value for money and price.
Risk	Technology and investment risks are high due to the factors above.	The technology risk is low, as technology and standards are in place and many vendors are available in the value chain; the investment risk is low as the market is now well understood.
Strategic focus	First to market and maintenance of leadership position technologically.	Cheapest vendor wins.

From the considerations above, maintaining an in-house development strategy can even be counter-productive: what made you the winner in the early adopter market can make you the loser in the mainstream market!

OUTSOURCE WHEN YOU CAN; DEVELOP WHEN YOU MUST!

The proposition of this white paper is that mainstream test and measurement equipment development and manufacturing should be based on an outsourcing model as much as possible to ensure that costs are kept to a minimum and internal resources are focused on the highest risk tasks.

A key premise is the availability of off-the-shelf technology that can be used to build high-performance test and measurement systems; perhaps not for leading edge development for the early adopter market, but definitely for mainstream markets.

This white paper will concentrate on one type of test and measurement system, an Ethernet Traffic Generator, as an example. Is it possible to build a high-performance Ethernet Traffic Generator based on off-the-shelf components?

The analysis will be based on a 10 Gbps Ethernet line rate, as this has now entered mainstream adoption. 40 Gbps Ethernet and 100 Gbps Ethernet are still early adopter markets, but with the ratification of key standards, 40 Gbps will be the next technology to "cross the chasm". Test and measurement equipment vendors should therefore begin planning now for how to make sure that this crossing is made successfully.

BUILDING AN AFFORDABLE 10 GBPS ETHERNET TRAFFIC GENERATOR

The keyword for success in the mainstream market is AFFORDABILITY. The goal for the test and measurement equipment vendor should be to make sure that every engineer who needs a test and measurement system can get one. The reality today is that these systems are so expensive that many development engineers cannot get access to these systems. Cost-conscious mainstream customers will only invest in test equipment for critical needs, such as test labs.

So, how can an affordable traffic generation system be built using off-the-shelf components and software?

Some of the key components required are:

- A hardware platform with adequate processing power and memory capacity
- A high-performance data transmission interface
- Ethernet frame generation application software

All of these components are available from commercial vendors:

- Standard servers now provide the processing power or open-source and memory capacity to handle high-speed, high-performance data communication applications.
- Intelligent network adapters are available that provide full line rate Ethernet frame transmission capability.
- Open-source software is available for generating Ethernet frames.

Successful demonstrations have been made of such an Ethernet traffic generation solution transmitting up to 20 Gbps of traffic based on off-the-shelf components of the type mentioned above.

The benefit of using these components is that it can save a considerable amount of development and manufacturing cost. Over 6 million tower and rack servers are sold per year[1] and chipset manufacturers are using this volume to invest in faster chips with more processing cores and faster memories. Performance of these chips is increasing at a rate of 60% per year at the time of writing. This provides a cost-performance combination, which is almost impossible to match over the long term with in-house development.

Standard Network Interface Cards (NICs) have, on the other hand, not proven to be capable of providing full line rate transmission. This is no surprise, as this is rarely called for in the main application of these technologies. To fill the gap, dedicated and focused vendors have created intelligent network adapters built specifically for applications that need to receive and transmit at full line rate. These solutions have also reached a volume level that allows a cost-performance combination, which makes it difficult to justify in-house development.

Using the combination of a standard server and intelligent network adapters allows any application software to be supported with high performance at full line rate, even up to 20 Gbps. This has allowed open-source software like Ostinato and PackETH to be successfully used for Ethernet traffic generation applications.

Let's take a closer look at the performance of these solutions.

STANDARD SERVER PLATFORMS

One of the advantages of the large PC market is that there is a constant development effort in new chips and architectures for both clients and servers. The dominance of Intel and AMD in the high-volume PC market provides a strong business case for continuous development effort. This has resulted in a constant improvement in PC server chipset performance culminating in the latest generation of chips and chipset architectures.
As an example, let's take a closer look at the Intel Nehalem architecture based on the Intel 5500 CPU.

[1] Source: Dell Oro, Feb 2010

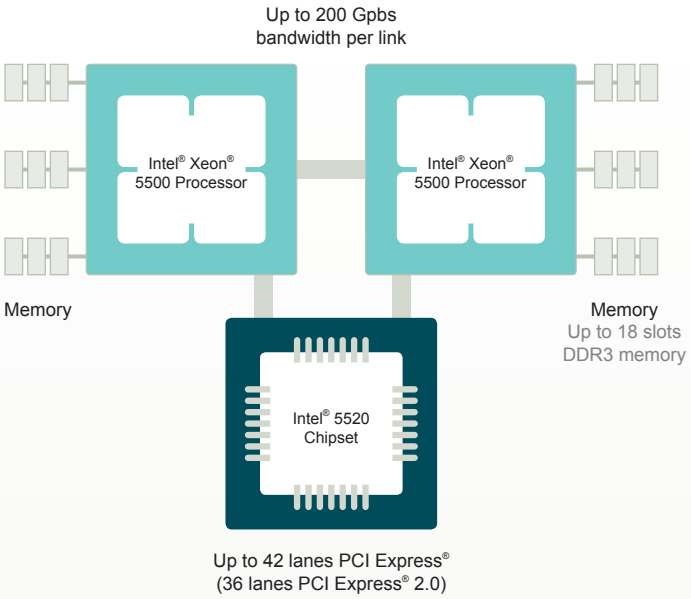

Figure 2: Intel Nehalem Architecture (Source: Intel data sheets, 2009)

According to Intel the Intel 5500 chipset provides up to 2.25 times performance improvement compared to previous Intel server chipset (the Intel 5400). One of the innovations of the Nehalem architecture is a more efficient memory architecture, which is one of the reasons why Intel can achieve the performance improvement quoted. The integration of the memory controller into the CPU chip itself avoids the use of a shared front-side bus and external memory controller, as was the case with the Intel 5400 chipset. The use of high-speed DDR3 RAM and a three-banked memory system provides memory bandwidth of more than 200 Gbps per CPU.

As memory access is often the limiting performance bottleneck of CPU processing, the introduction of a 3-6 times faster memory architecture combined with the introduction of a large layer 3 CPU cache and Symmetric Multi-Threading (SMT), referred to as Hyper Threading by Intel, have succeeded in addressing this bottleneck.

The introduction of Quick Paths (the high-speed links between the CPUs and the Controller chip) and a new very fast PCI Express controller solves two other performance bottlenecks so that data can now be transferred between network adapters and server memory and between CPU cores at rates of more than 100 Gbps.

From this short overview, it should be clear that the Nehalem architecture and equivalent chipsets have the processing resources and memory architecture to support even the most demanding real-time 10 Gbps applications. What's more, developments continue with servers now supporting 2, 4 and 8 CPU chips each with 4 to 8 cores per CPU chip.

NETWORK ADAPTERS

Network Interface Cards (NICs) are familiar to anyone with a PC server or client. These are the interface cards that provide Ethernet interfaces for connecting to the Local Area Network (LAN) or the Internet.

This communication task is also what NICs were designed to perform, namely sending specific Ethernet frames[2] to a specific destination denoted by the Ethernet Media Access Control (MAC) address. In these cases, the NIC need only concern itself with the Ethernet frames that are directly addressed to the MAC addresses associated with the ports on the NIC. All other traffic can be ignored.

NICs designed for server applications (server NICs) are optimally designed to handle large frames. The reason for this is that most server applications use the TCP protocol for transfer of data. TCP will seek to use as few and as large frames as possible to transfer data. Using large frames has the advantage for the server application that it reduces protocol overhead, which is directly proportional to the number of frames to be handled. In other words, larger frames means fewer frames with less overhead and more payload throughput.

FULL LINE RATE TRANSMISSION

Transmission of 10 Gbps data per port at full line rate is a challenge for server NICs. For traffic generation, this is an absolute requirement. Not only that, but it is a requirement that this performance can be met for all Ethernet frame sizes.

[2] The term "packets" is usually reserved for layer 3 protocols such as the Internet Protocol (IP). For layer 2 protocols, such as Ethernet, the header and encapsulated data information is collectively referred to as a "frame". It is therefore not strictly correct to refer to an "Ethernet packet".

Since server NICs are built to handle large frames, transmission of these large frames can be achieved at a rate close to the theoretical maximum. However, as frame sizes get smaller, server NICs can no longer keep up:

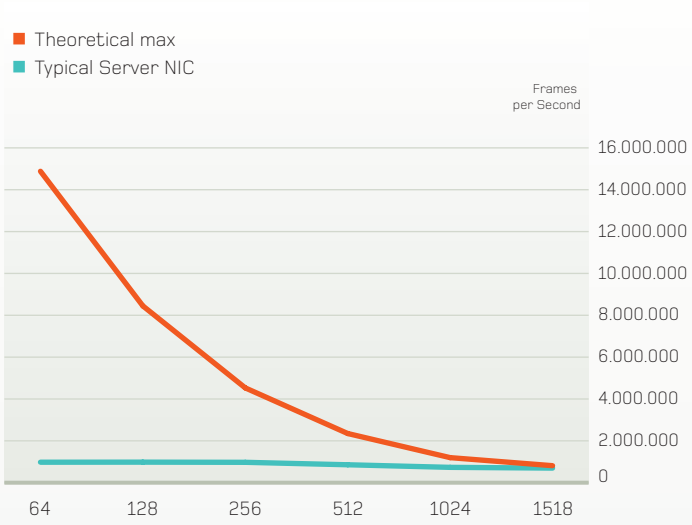

Figure 3: Ethernet Frames Per Second Transmission per 10G port

As can be seen, nearly 15 million Ethernet frames per second need to be transmitted per second when the Ethernet frame size is 64 bytes. This is far above what a typical server NIC can achieve. This is a real use case for Ethernet Traffic Generation solutions, which would be expected of being able to provide the theoretical max transmission rate as shown above.

Fortunately, intelligent network adapter vendors have designed their products to provide exactly this type of performance. For example, Napatech's NT20E2 intelligent network adapter can provide sustained transmission at theoretical max for all frame sizes on 2 x 10 Gbps ports, matching the requirements above.

In addition, the NT20E2 provides precise control of transmission timing via configuration of the Inter-Frame Gap (IFG) and automatic generation of Ethernet, IP, TCP and UDP checksums.

This off-loads and accelerates the Ethernet traffic generation application software providing extra processing power for more advanced features.

OPEN-SOURCE TRAFFIC GENERATION TOOLS

A number of open-source traffic generation tools are now available and can be implemented directly on a standard server with intelligent network adapters. A complete list will not be provided here, but two popular open-source applications will be mentioned as examples.

OSTINATO

Ostinato is a cross-platform, open-source traffic generator software application. It provides a user-friendly GUI where it is possible to configure traffic streams based on various popular communication protocols.

For more information see: http://code.google.com/p/ostinato/

PACKETH

PackETH is similar to Ostinato, but focused on Ethernet and Linux. It also provides a user-friendly GUI where it is possible to configure traffic streams including protocol fields for communication protocols typically supported by Ethernet.

For more information see: http://packeth.sourceforge.net/

SUCCESSFUL 20 GBPS ETHERNET TRAFFIC GENERATION DEMONSTRATION

Napatech has successfully demonstrated a 20 Gbps Ethernet traffic generation solution based on a standard server, Napatech's NT20E2 intelligent network adapter and Ostinato application software. This demonstration was recently shown at the Interop New York event in October 2010.

For a video of the demonstration see:
http://www.napatech.com/resources/videos.html

CHALLENGE FOR TEST AND MEASUREMENT EQUIPMENT VENDORS

As shown above, it is now possible to build high-performance 20 Gbps traffic generation solutions based on off-the-shelf components. This provides a challenge to test and measurement equipment vendors, as it is now possible for potential customers to build their own solutions. Competition is no longer from other test and measurement equipment vendors, who are facing similar challenges, but in-house development of test solutions.

If test and measurement equipment vendors do not address this challenge, they can find themselves confined to early adopter markets and high-end opportunities. As off-the-shelf systems improve in performance, they can gradually undermine high-end opportunities leaving test and measurement vendors with the high-risk, lower-volume early adopter market as their only market opportunity.

Alternatively, test and measurement equipment vendors can re-think their strategies and realize that a different approach is needed for addressing the needs of the mainstream market; an approach, which re-focuses development resources on value-adding software development by taking advantage of the same off-the-shelf hardware components that end users could potentially take advantage of for in-house development. This allows test and measurement equipment vendors to not only protect their current market, but even expand to making test and measurement equipment affordable for all engineers in the organization.

ABOUT NAPATECH

Napatech is the leading OEM supplier of multi-port 10 GbE and 1 GbE intelligent adapters for real-time network analysis with over 70,000 Ethernet ports deployed. Napatech network adapters provide real-time packet capture and transmission with full line rate throughput and zero packet loss no matter the packet size. Intelligent features enable off-load of data traffic processing and packet analysis normally performed in the CPU. This results in more processing power for the network monitoring, analysis, management, test, measurement, security or optimization application being supported. Napatech has sales, marketing and R&D offices in Mountain View (CA), Andover (MA), Washington D.C., Tokyo (Japan), and Copenhagen (Denmark).

Europe, Middle East and Africa
Napatech A/S
Tobaksvejen 23 A, 1
DK-2860 Soeborg
Denmark

Tel. +45 4596 1500
Fax +45 6980 2970
www.napatech.com
nteusales@napatech.com

Americas
Napatech, Inc.
One Tech Drive
Suite 110
Andover, MA 01810
US

Tel. +1 888 318 8288
Fax +1 978 824 9414
www.napatech.com
ntamericassales@napatech.com

APAC
Napatech
1/F Place Canada
7-3-37 Akasaka Minato-ku
107-0052 Tokyo
Japan

Tel. +81 3 6894 7678
Fax +81 3 6894 7701
www.napatech.com
ntapacsales@napatech.com