



Technical Brief

NVIDIA DualNet with Teaming Advanced Networking

May 2006
TB-02499-001_v01

Table of Contents

NVIDIA DualNet with Teaming	4
What is NVIDIA DualNet?.....	4
Teaming	6
Intelligent Transmit and Receive Load Balancing.....	7
Fail-Over	7
TCP/IP Acceleration	8
Home Gateway.....	10
Applications of NVIDIA DualNet	11
Teaming Scenario	11
Load Balancing and Fail-Over Scenario.....	11
Conclusion	11

NVIDIA DualNet with Teaming

Today's desktop enthusiast-level computers are used for more than just a gaming platform., they also double as game servers at LAN parties, delivering not only the game itself to a number of connected computers, but also supplying larger files such as game updates and levels to these same computers. When sharing large amounts of data over a network, it is important that the computer serving the files not be starved of network bandwidth, especially if that computer is also being used as a gaming client.

The NVIDIA® DualNet™ technology brings a number of advanced enterprise-level network technologies to the enthusiast including teaming and TCP/IP acceleration. Teaming doubles the network link by combining the two integrated Gigabit Ethernet ports into a single 2-Gigabit Ethernet connection to increase the overall link speed with high availability and redundancy. TCP/IP acceleration reduces CPU utilization by offloading CPU-intensive packet processing tasks to hardware using a dedicated processor for accelerating traffic processing combined with hardware-optimized software.

What is NVIDIA DualNet?

NVIDIA DualNet technology integrates dual Gigabit Ethernet MACs that, by design, eliminate network bottlenecks and improve overall system efficiency and performance, resulting in the industry's fastest Gigabit Ethernet. These two Gigabit Ethernet ports can be used individually, combined, or configured in a number of ways depending on the needs of the user.

A number of motherboards have offered dual-Gigabit Ethernet ports, however, all of these offerings do so using two different chips. The NVIDIA nForce® MCP is the first integrated motherboard solution to offer dual-Gigabit Ethernet using the same physical chip (Figure 1).

DualNet is more than just the combination of two gigabit ports in hardware. DualNet allows a PC to serve as a home gateway and it provides advanced networking features including teaming, load balancing, fail-over, and TCP/IP acceleration.

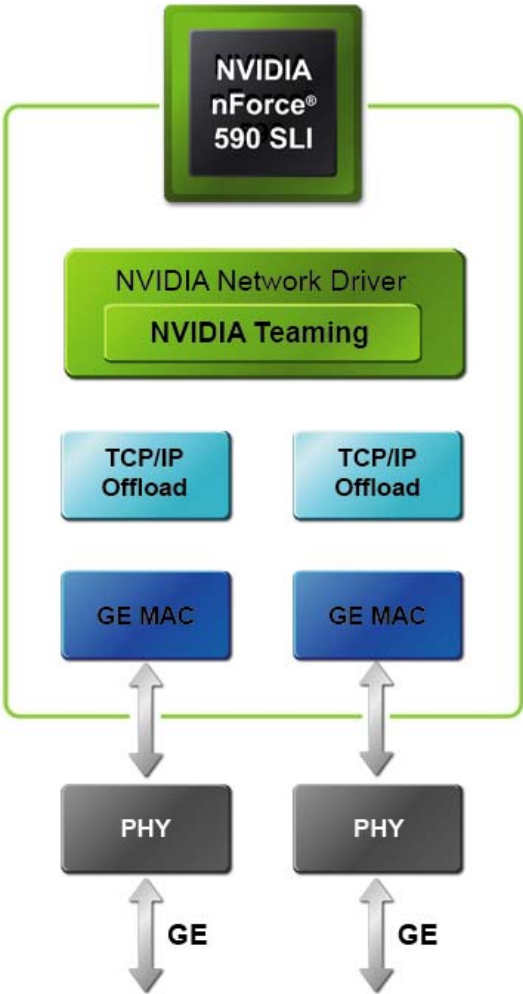


Figure 1. Dual-Gigabit Ethernet using One Physical Chip

Teaming

Teaming allows both of the Gigabit Ethernet ports in NVIDIA DualNet configurations to be used in parallel to increase the overall link speed of the Ethernet connection.

NVIDIA DualNet with Teaming is an easy way to set up a high-speed, 2-Gigabit Ethernet backbone that allows multiple computers to be connected simultaneously at full gigabit speeds to unleash the power of your PC (Figure 2).

When Teaming is enabled, network performance can be increased by connecting multiple gigabit ports, and intelligently balancing traffic between them. Integrated fail-over technology ensures that if one link goes down, traffic is instantly and automatically redirected to the remaining link. This in turn provides all the capabilities that are needed to implement a highly available high performance file server.

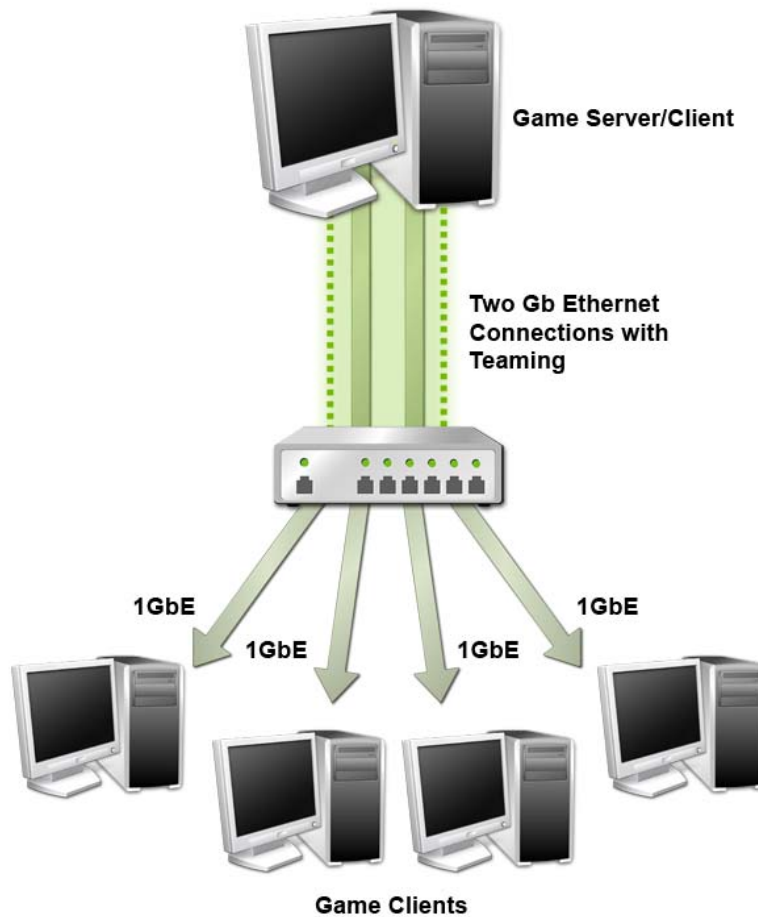


Figure 2. NVIDIA DualNet with Teaming

Intelligent Transmit and Receive Load Balancing

When Teaming is enabled, multiple gigabit links within the team maintain their own dedicated MAC address, but the combined team shares a single IP address. Transmit load balancing uses the destination (client) IP address to assign outbound traffic to a particular gigabit connection within a team. When data transmission is required, the network driver uses this assignment to determine which gigabit connection will act as the transmission medium. This ensures that all connections are balanced across all the gigabit links in the team. If at any point one of the links is not being used, the algorithm dynamically adjusts the connection to ensure optimal connection utilization.

Receive load balancing uses a connection steering method to distribute inbound traffic between the two gigabit links in the team. When the gigabit ports are connected to different servers, the inbound traffic from different servers is distributed between the links in the team.

Fail-Over

Fail-over is the capability of the NVIDIA nForce MCP to automatically switch over to a redundant or standby connection should one of the gigabit ports fail, drop out, or is disconnected (Figure 3). Fail-over is typically required in server configurations to increase server reliability and to ensure continuous availability.

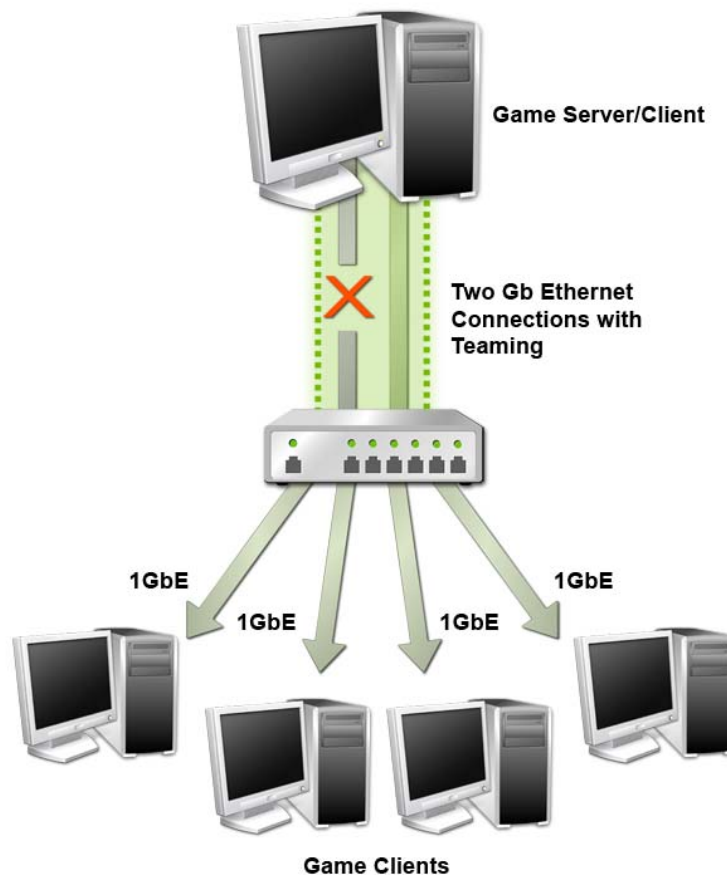


Figure 3. Fail-over Configuration

Fail-over works closely with load balancing to ensure that network traffic is rerouted properly if a link fails. For example, if a cable were unplugged from one gigabit port, the network traffic is diverted to the remaining gigabit link. This is an automatic process that requires no user intervention. Once the lost link has been restored, the grouping is reestablished and traffic begins to flow on the restored link.

TCP/IP Acceleration

NVIDIA nForce MCPs have TCP/IP acceleration and hardware offload capability built in to the native Gigabit Ethernet Controller. This lowers the CPU utilization when processing network traffic at gigabit speeds. The hardware offload is fully compliant with Microsoft's Chimney architecture.

NVIDIA TCP/IP Acceleration is a networking solution that includes both a dedicated processor for accelerating networking traffic processing and hardware-optimized software. TCP/IP acceleration provides deep levels of networking and traffic inspections at full-duplex Gigabit Ethernet speeds. By offloading CPU-intensive packet processing tasks in hardware, TCP/IP acceleration delivers the highest system performance (Figure 4).

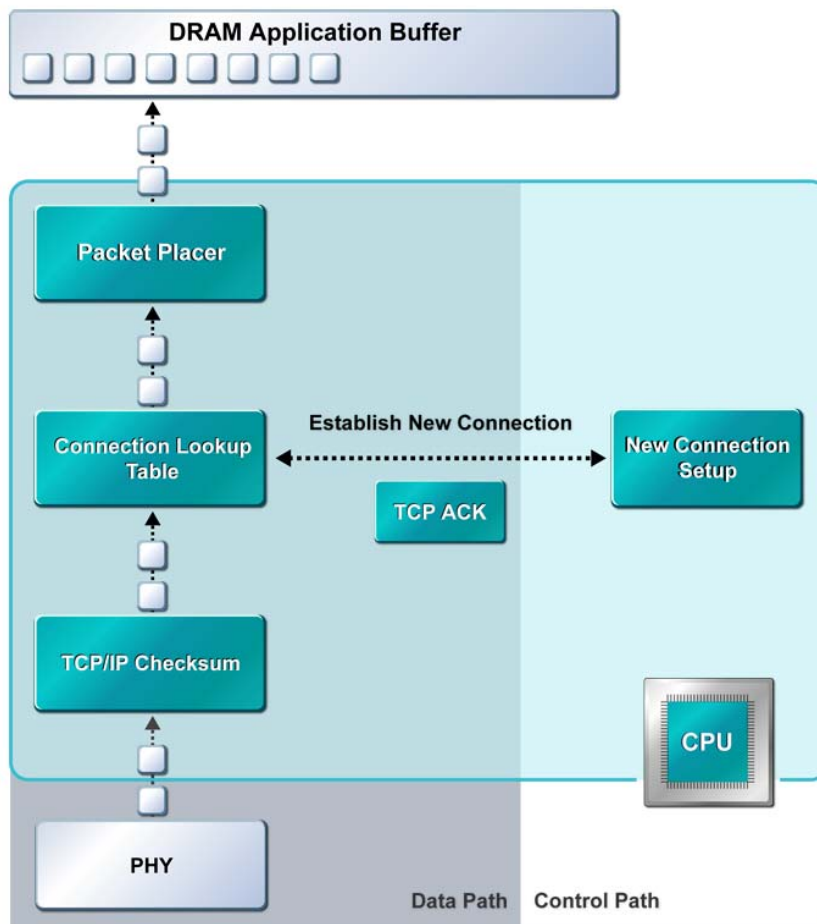


Figure 4. TCP/IP Acceleration

In traditional networking environments, inspecting packets is laborious and affects CPU overhead, memory bandwidth, and overall system latency. For example, in pure software solutions, the CPU is responsible for processing all aspects of the TCP protocol: Checksumming, ACK processing, and connection lookup. Combined, this puts a significant load on the CPU since the checksumming alone requires the CPU to inspect every data byte of the incoming stream.

By comparison, NVIDIA TCP/IP acceleration (Figure 5) processes network packets in hardware, allowing them to use an “express lane” that bypasses the traditional CPU bottleneck, thereby improving overall throughput and lowering CPU utilization.

Similar to the prior example, all packet data is processed and checksummed, but instead of forcing data to be moved to the CPU for software-based processing, the embedded hardware engines in the MCP take over. Each data byte can be processed inside the MCP, thereby offloading a fixed computational task from the CPU which is then free to continue normal operation on other required tasks to keep PCs responsive, even during large file downloads.

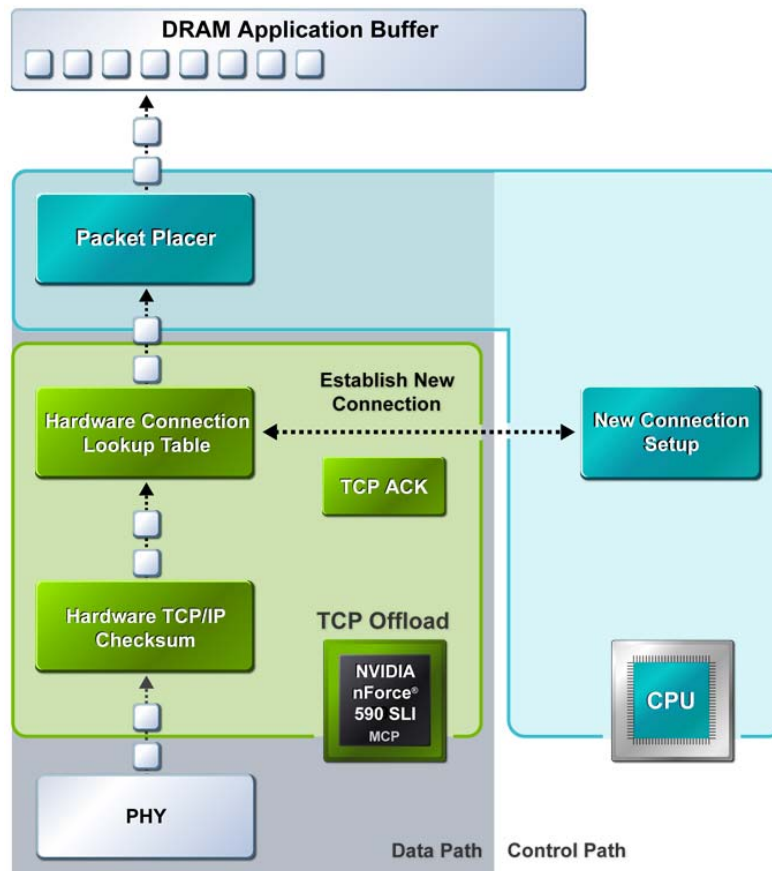


Figure 5. NVIDIA TCP/IP Acceleration

NVIDIA's TCP/IP Acceleration Technology is enabled using the NVIDIA control panel. If a software firewall is installed on the system, enabling TCP/IP Acceleration technology may cause some network traffic to bypass the firewall. A warning message is displayed explaining this if a firewall is present and a user enables TCP/IP acceleration.

Home Gateway

NVIDIA DualNet technology allows a PC to be used as a home gateway. Using the two integrated Gigabit Ethernet ports, network traffic can be expanded and managed between multiple clients.

For example, data from a wide area connection (WAN) such as a cable or DSL modem that delivers Internet traffic can be attached to the first Gigabit Ethernet connection on an NVIDIA nForce platform while the second Gigabit Ethernet connection can be connected to another client or local area connection (LAN) comprised of a network hub or switch. The intelligence behind NVIDIA DualNet technology ensures that the rapid transfer of data is handled without any added arbitration or latency (Figure 6).

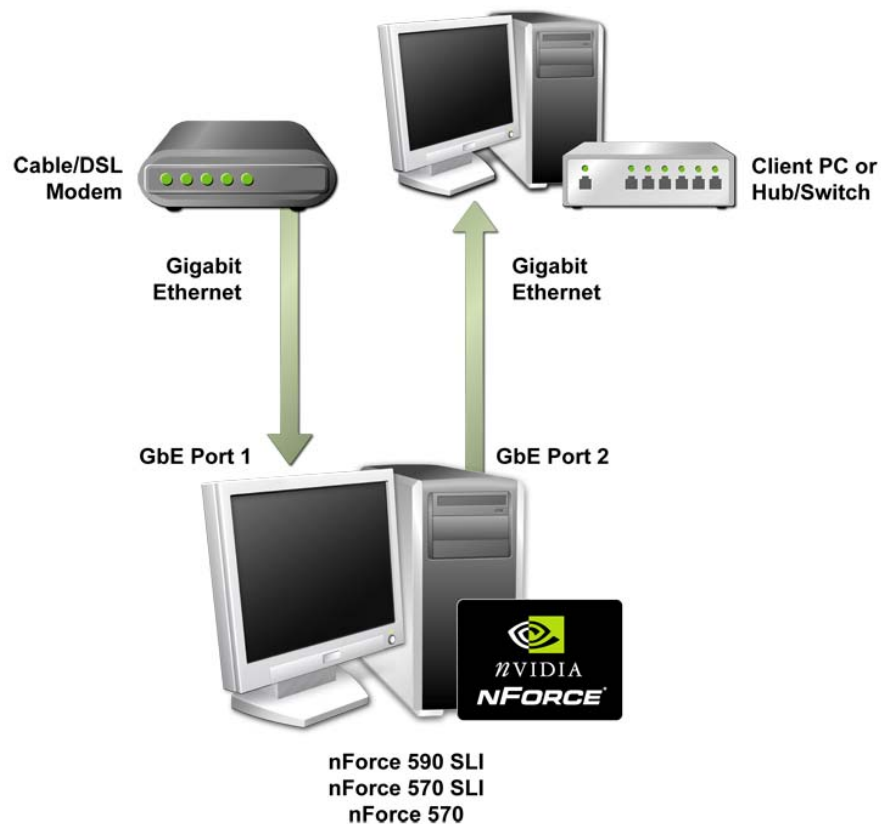


Figure 6. NVIDIA DualNet Technology in a Home Gateway Configuration

Applications of NVIDIA DualNet

Teaming Scenario

In a LAN gaming environment, often held at LAN parties, enthusiasts physically bring their PCs together to play games. Often, one of person hosts the game, turning their PC into a game server. That person not only hosts the game, but often provides large game files in the form of patches and levels to the rest of the group. As people arrive, they copy these files and prepare their PC to join the game. If that server is already hosting a game, the bandwidth available to host the clients running the game can be compromised and can actually affect game play.

One way to ensure bandwidth is to enable NVIDIA DualNet with teaming on the computer acting as the game server. When teaming is enabled, the two single Gigabit Ethernet ports are combined into a single 2-gigabit connection, doubling network bandwidth to game clients and reducing the impact of multiple downloads to the gaming clients.

Load Balancing and Fail-Over Scenario

If both gigabit ports on the motherboard are connected to the same switch, load balancing is enabled using a single IP address, and one of the Ethernet cables is unplugged, the network connection is retained. If this happens while streaming Internet video, the video stream will continue without dropping frames. If a file is being downloaded, the download will continue without loss of packet or corruption of data.

Conclusion

NVIDIA nForce MCPs not only integrate dual Gigabit Ethernet MACs that by design eliminate network bottlenecks and improve overall system efficiency and performance, they include a number of enterprise-level network technologies to enthusiasts.

TCP/IP Acceleration dramatically lowers CPU utilization when processing network traffic at gigabit speeds and Teaming allows both of the Gigabit Ethernet ports in NVIDIA DualNet configurations to be used in parallel to increase the overall link speed of the Ethernet connection. The load balancing and fail-over technologies deliver increased performance and reliability by intelligently assigning connections across multiple Ethernet ports.

NVIDIA DualNet technology delivers the highest network throughput at the lowest CPU utilization. The manageable and stable NVIDIA networking solution results in better networking management and a lower total cost of ownership. Only NVIDIA integrates this level of networking features to allow you to take your online experience to the next level.



Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, nForce, and DualNet are trademarks or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated

Copyright

© 2006 NVIDIA Corporation. All rights reserved.



NVIDIA.

NVIDIA Corporation
2701 San Tomas Expressway
Santa Clara, CA 95050
www.nvidia.com