



Technical Brief

**NVIDIA nForce3 Professional
Media and Communication
Processors**
Enterprise-Class Networking



Enterprise-Class Networking for Today's Professionals

Workstation users demand excellent overall system performance and a lot more. The ability to get the job done often involves accessing a network—network-based applications, sending/receiving e-mail and files, streaming media, and handling large and complex data sets over a network are commonplace tasks. While CPU and graphics processing unit (GPU) speeds and system capacities advance at an aggressive pace, the basic system architectures evolve more slowly, and critical subsystems for interfacing to the outside world and handling data can often limit the overall desktop experience. The NVIDIA nForce™ core technology solved this problem by delivering advanced core technologies and industry-leading performance to PC users. Today, the NVIDIA nForce3 Professional media and communication processors (MCPs) solve the same problems for workstation users.

The newest NVIDIA nForce3 Professional MCPs include third-generation NVIDIA networking technology. The Company's proven expertise with streamlining core system functionality has yielded a strong suite of built-in enterprise-class networking solutions spanning hardware, drivers, and user interface functionality. The full-featured NVIDIA networking technology includes capabilities that address:

- ❑ **Advanced access**—NVIDIA MCP technology provides a choice of popular and emerging high-speed Ethernet technologies.
- ❑ **Management**—monitoring critical system events and delivering early-warning alerts ensures that network managers are notified before problems escalate and impact critical applications and users.
- ❑ **Security**—built-in, streamlined functionality enabling the protection of systems from unauthorized access benefits every user on the network.

Optimized to take advantage of today's system architectures, the NVIDIA enterprise-class networking solution offloads many compute-intensive networking operations from the CPU, contributing to increased overall workstation performance.

Advanced Access

The NVIDIA nForce3 Professional MCPs deliver field-proven third-generation NVIDIA Media Access Control (MAC) technology. This commercial-grade solution gives system designers an industry-standard Media Independent Interface (MII) or Reduced Gigabit MII (RGMI) for 10/100BASE-T Fast Ethernet and 1000BASE-T Gigabit Ethernet. Gigabit Ethernet meets the high-performance requirements for commercial users on today's high-speed backbone networks. Adherence to the industry standards means that system designers can interface the NVIDIA solution to Ethernet PHYs from a variety of vendors.

The integrated Ethernet solution offered with the NVIDIA nForce3 Professional solutions ensures an optimized design with appropriate hooks into all of the system hardware and NVIDIA software. Compared to external Ethernet offerings, the NVIDIA solution ensures that future upgrades and enhancements incorporate changes to the appropriate part of the complete data path.

The NVIDIA implementation of task offloads incorporates support for the latest Internet Protocol (IP) specification, IPv6. This forward-looking design ensures a forward-compatible system and offers system designers investment protection as the standards evolve.

Enhanced System Design

By integrating Fast Ethernet and Gigabit Ethernet technology on a single-chip solution, NVIDIA nForce3 Professional MCPs offload the CPU and increase overall system performance. The specific features for task offloading include:

- ❑ **Checksum offload:** calculating checksums is the most CPU-intensive function in the networking stack, contributing to long path length and causing cache churning. Offloading this to the NVIDIA nForce3 Professional networking solution increases system performance, improves CPU cache effectiveness, and raises overall network responsiveness—even for systems where the CPU is not otherwise busy.
- ❑ **TCP segmentation (large send) offload:** for large send operations, the NVIDIA nForce3 Professional offloads the CPU from the work required to segment a large buffer into packets that fit within the network.
- ❑ **Jumbo frame support:** larger maximum transmission units (MTUs) or jumbo frames reduce the per-byte overhead for large transfers. Offloading the handling of jumbo frames to the networking platform reduces the number of calls to the network driver, thereby reducing CPU overhead.
- ❑ **Interrupt moderation:** the NVIDIA nForce3 Professional networking solution can process several packets with only one CPU interrupt, instead of generating one interrupt per packet. This increases overall networking performance by minimizing interrupt latency and cache churning.

In addition to these performance enhancing features, the NVIDIA nForce3 Professional platforms give system designers other benefits:

- ❑ **Reduced latency** when compared to external Ethernet solutions
- ❑ **Increased Gigabit Ethernet throughput** (the on-chip design uses high-speed data busses, and allows this interface to exceed the speed of PCI-based solutions)
- ❑ **Increased PCI performance** (PCI Gigabit Ethernet solutions quickly saturate the bus and degrade PCI component performance, while the NVIDIA nForce3 Professional Gigabit Ethernet technology resides on an internal high-speed bus and, therefore, does not impact the performance of PCI devices)
- ❑ **Lowered power and heat** when compared to external Ethernet solutions
- ❑ **Increased reliability and testability** resulting in higher quality

Benefits to the User and Network Manager

The NVIDIA nForce3 Professional on-chip Ethernet solutions offer:

- ❑ **Configuration flexibility:** for interoperability with low-end switches, the auto-negotiation plus feature allows the Fast Ethernet mode to be set from the client side when connected to a non-managed switch. (Gigabit Ethernet incorporates a different auto negotiation feature tailored to the specific operation of gigabit transfers.)
- ❑ **Flow control:** for both Fast and Gigabit Ethernet transfers, the transfer speed is adjusted depending on network conditions.
- ❑ **Traffic prioritization:** the NVIDIA driver supports the IEEE 802.1p traffic prioritization specification, and the IEEE 802.1Q virtual local area network (VLAN) specification. IEEE 802.1p allows each frame to be labeled with one of eight priority codes, which results in optimized transmission through intelligent networks.
- ❑ **Power management:** the NVIDIA solution complies with the Advanced Configuration and Power Interface (ACPI) 2.0 standard. ACPI establishes open industry-standard interfaces for OS-directed configuration and power management on laptops, desktops, and servers.
- ❑ **Remote wake-up:** this feature enables receiving a remote wake-up packet so that a network manager can access the system even if it had been previously put to sleep or shut down. (Full functionality of this capability requires OS support.)
- ❑ **Network boot:** the NVIDIA nForce3 Professional MCPs implement the Preboot Execution Environment (PXE) standard, allowing enterprise network clients to download software images and configuration parameters from a centralized server before the operating system is loaded. This essentially allows a network interface to function as a boot device, letting network managers remotely set up new systems, perform pre-OS system management, and remotely boot over the network.
- ❑ **Network management:** a full suite of interfaces—command line interface (CLI), Web browser interface, and WMI scripts—gives network managers an option that fits into their environment.

Management

To facilitate effective network management practices, the NVIDIA networking solution complies with the Alert Standard Format (ASF) 1.03 specification. Alerts are generated for a variety of events, resulting in reduced maintenance costs and an increase in uptime and customer satisfaction. ASF enables the generation of these alerts even if the system is powered off or the operating system is not yet loaded. Once a system reports a warning or error, ASF allows corrective action to be taken, including remote reset, power-on, or power-off functions. Many system problems can be caught early before there is any negative impact on users. The generated alerts fall into the categories of remote monitoring (thermal and cooling system status), anti-theft measures (case intrusion and CPU power status), and operating system monitoring (boot failures). Tables 1 and 2 list all of the alerts supported.

NVIDIA nForce3 Professional support for ASF alerts can be used in conjunction with Simple Network Management Protocol (SNMP) applications that recognize ASF event notification. For networks that do not use any of these applications, NVIDIA offers a monitoring program that can recognize the alerts and notify the user or network manager.

Table 1. NVIDIA nForce3 Professional: Supported ASF 1.03 Internal Event Alerts

Alert	Description (event that triggers the alert)
Total Cost of Ownership (TCO) timeout	System fails to boot
Fan problem	Rotation speed is too low (could result in increased CPU or system temperature)
System Management Bus (SMBus) alert	SMBus slave address matching condition
Wake on LAN (WoL) or Alert on LAN	Remote wake-up initiated
Internal ASF timer/heartbeat	At programmed intervals, this heartbeat signal is generated (indicating that the system is up and running)
OS hung	The operating system is hung, and device interrupts are not being serviced
System on/off or wakeup (PWRBTN#)	Each time the power button is pressed for turning the system on or off (typically used to detect someone turning off the system during an attempted theft)

Table 2. NVIDIA nForce3 Professional: Supported ASF 1.03
Generic General-Purpose Input/Output (GPIO) Events

Alert	Description (event that triggers the alert)
Thermal alarm (THERM#)	Temperature in the system exceeds acceptable limits
CPU hot (PROCHOT#)	CPU temperature is increasing
CPU overheated (THERMTRIP#)	CPU temperature exceeds threshold (possible fan failure; requires system shutdown)
Case open (INTRUDER#)	Case intrusion (possible theft attempt)
SMBus alert pin (SMBALERT#)	Configurable (tied to SMBus)
Generic system-defined GPIO alert	Configurable (tied to GPIO)

Security

Today's networks require extensive measures to protect systems and applications from unauthorized access and malicious acts that can result in lost data, network outages, downed systems, and other disasters that can directly impact a company's bottom line.

A high-performance, network-level firewall, NVIDIA Firewall protects your system from intruders by filtering unauthorized traffic. Integrated into NVIDIA nForce3 MCPs with NVIDIA Gigabit Ethernet, it provides professional-grade traffic inspection capabilities, advanced management features--remote access, configuration, and monitoring--and is easy to use and setup via a user friendly wizard.

NVIDIA Firewall

Firewall functionality provides protection from unauthorized users attempting to break into private data repositories or attempting more malicious behaviors such as overwhelming a server or an entire network. Companies typically establish a firewall that limits access from non-employees, but there is often no protection from attempted break-ins from within the organization—an increasingly common security problem. Traditional firewalls also offer no protection for off-site employees working from the road or home. NVIDIA security technology meets all of these needs with another level of defense at each end-point (i.e., desktop). The NVIDIA Firewall technology protects each system from attempted breaches, or break-ins that can happen when a user is connected from a remote site using an “always-on” broadband link to the Internet.

NVIDIA Firewall technology provides support for stateful and stateless firewall inspection. The NVIDIA approach is hardware optimized to provide excellent protection and throughput with minimal CPU utilization. The NVIDIA Firewall

supports features formerly available only in high-end firewall devices. The NVIDIA nForce3 Professional MCPs bring a high-performance, no-compromise professional firewall to workstations as integrated with Gigabit Ethernet MAC driver functionality.

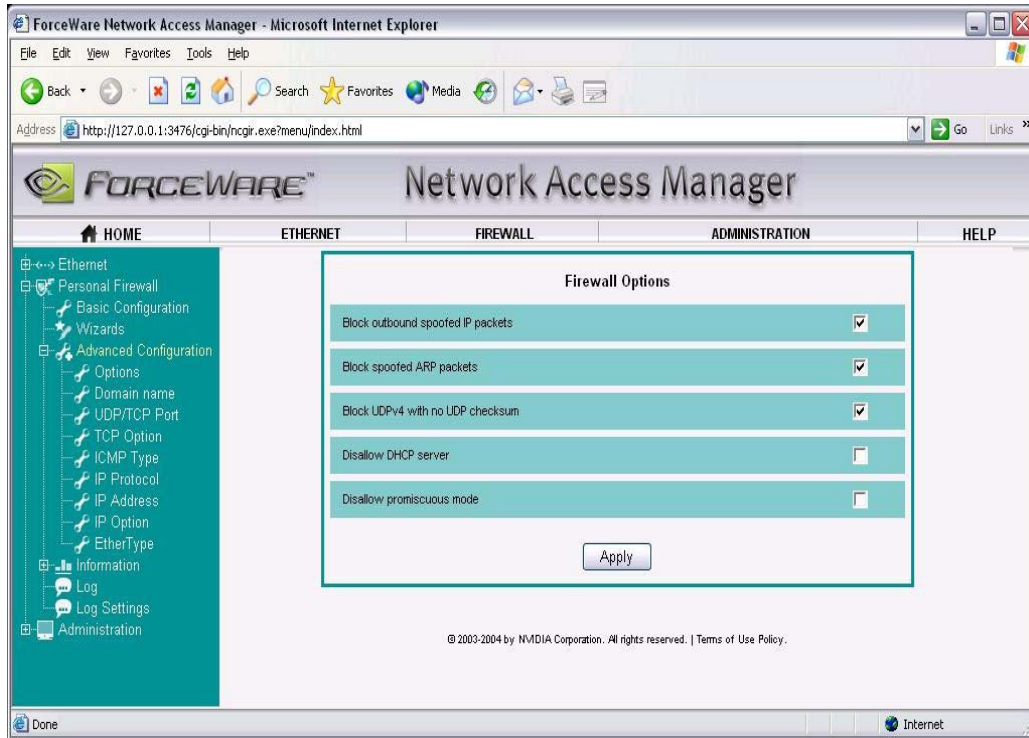


Figure 1. The NVIDIA Firewall can be easily configured with a Web-based browser interface.

Anti-hacker Features

The NVIDIA nForce3 Professional MCPs offer protection from “IP spoofing,” a common practice used by hackers to mask the source of malicious traffic. The NVIDIA security technology includes checks for inconsistent IP addresses, and eliminates this type of traffic at the source. In addition to anti-spoofing capabilities, the NVIDIA nForce3 Professional MCPs include features for enabling protection from anti-sniffing (stop listening to traffic not address to the desktop), anti-ARP cache poisoning, and prevent system from acting as a DHCP server (for example, stop assigning illegal IP addresses).

In a corporate setting, an end-point firewall (i.e., desktop firewall) with anti-hacking capabilities can serve to reduce the internally originated security breaches, and can inhibit desktops from generating unauthorized traffic hence improving the overall security while reducing the IT staff resource requirements. (See Figure 1.)

Conclusion

The NVIDIA nForce3 Professional enterprise-class networking technology gives both network managers and workstation users the benefits of in-depth NVIDIA experience for optimizing critical core functionality. With built-in solutions for today's fastest Ethernet access methods, the NVIDIA nForce3 Professional platform processors ensure that systems can survive and thrive in today's fast evolving networked environments. Management functionality has been built into the NVIDIA solution, making it possible to effectively monitor and catch system problems before they escalate. The NVIDIA solution also incorporates high-performance technologies for security features, introducing a desktop level of protection that augments other corporate security measures and makes in-home and corporate systems safer from unauthorized access.



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.

NVIDIA and the NVIDIA logo are registered trademarks and NVIDIA nForce is a trademark of NVIDIA Corporation.

Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright NVIDIA Corporation 2004.



NVIDIA.

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