



Using Efficient and Cost-Effective Connectors In Performance and Power Optimized Applications

A Tehuti Networks white paper

Blaine Kohl, VP of Marketing

How to contact us

For inquiries regarding Tehuti Networks, please use the information below.

Email:

info@tehutinetworks.net

Website:

www.tehutinetworks.net

THIS DOCUMENT IS PROVIDED BY TEHUTI NETWORKS.

Tehuti Networks, the pyramid logo, and OptiStrata are trademarks of Tehuti Networks Ltd. All other trademarks are the property of their respective owners. Tehuti Networks Ltd. may change the specifications and product descriptions at any time, without any notice. No license, express or implied, to any intellectual property rights is granted by this document. While every attempt has been made to assure that the information presented in this document is accurate, Tehuti Networks assumes no liability whatsoever relating to fitness for a particular purpose, merchantability or infringement of any patent, copyright or other intellectual property rights.

Version 1.1.

January 2007

© Copyright 2007 Tehuti Networks Ltd.

Executive Summary

To meet the increasing needs of server, storage, and network appliance workloads, equipment manufacturers are looking for ways to increase packet processing to improve system performance. However, for a large percentage of the market, many other parameters must be considered in addition to performance. Power and thermal issues are typically cited as top concerns of data center administrators and IT professionals, both in terms of availability of power and cost of cooling. In addition, the data center real estate is expensive. This has translated into a desire to utilize highly-integrated, modular systems to minimize the footprint and number of components on boards. With increased density and thermal conditions, connectors can become one of the most significant system performance impairments. This article addresses alternative ways equipment manufacturers can address these critical parameters.

Networking Adoption Trends

The industry has transitioned to Gigabit Ethernet (GbE) speeds and is now shifting to 10 GbE. Increasingly, more expensive and higher performance processors, connectors and cable or PCB traces are required to keep pace with network throughput. Due to design constraints, the memory subsystem cannot physically keep pace with network and the Shannon capacity of cabling (connectors and cable) is taxing the capabilities of today's silicon. This means that latency is increasing to data-starved, waiting applications such as 10 GbE PHYs to compensate for channel impairments and processors that wait through hundreds of idle cycles while memory works at its slower speed.

Servers and networks are the cornerstone upon which today's industries operate. They perform functions as simple as email and web services to as complex as intrusion detection, sensor networks and grid computing. Nearly all of the data required for those applications moves between systems in the form of packets. The basic problem is that packet processing has to be accelerated, both by the network interface and by the host system, to prevent the system from being bogged down. However, these bottlenecks need to be solved in a way that are sensitive to power and footprint constraints.

This disparity, between performance and power, is impacting the performance of today's systems designed to handle the critical real-time transactions. With increasing demands presented by powerful applications like e-Commerce, medical imaging and data warehousing, today's servers are being designed to process larger files and move more data faster than ever before. This directly impacts the network interface and the host system.

Industry Responses

In many ways 10GbE is a straightforward extension of previous The ASIC industry has attempted to address the host system bottlenecks through a number of different approaches. Unfortunately, each has significant limitations that inhibit widespread adoption in future system designs.

Traditional approaches such as selective, albeit limited, software offloading of specific tasks

(for example, IP checksum) have been implemented for several years. However, as traffic volume and network bandwidth increase, software can provide very little performance gain.

TCP/IP offload engines (TOEs) try to address network I/O bottlenecks by assuming some of the burden of TCP/IP processing and freeing up host CPU cycles for applications. TOEs have been slow to materialize for 10 GbE due to the increased power, footprint and cost.

The RDMA protocol and an extension to TCP/IP, iWARP, allows one computer to transfer data directly into a specific destination memory on another computer. It minimizes demands on bandwidth, reduces the overhead for packet processing, and reduces latency. RDMA exists for 10 GbE, but similar to TOE, it suffers from increased power, footprint and cost.

Because of the higher power budgets associated with these types of solutions, the market has been hesitant to embrace these solutions where power and footprint constraints are critical. In addition, there is often significant rework of the protocol stack that must be done which increases integration difficulties. The increased power of these solutions has an impact on the supporting components. Connectors would need to handle the increased thermal gradients.

Because of the concerns regarding power and footprint, a different approach was introduced to address the limitations associated with TCP/IP processing to meet the above customers' requirements. Tehuti Networks discovered that a series of relatively minor, but important, changes could be made to enhance the performance without adversely affecting the cost, foot print, or the power consumption and dissipation. This approach is called host hardware offload.

Another means that system designers have at their disposal is better components. As mentioned earlier, connectors are often the largest source of impairments due to impedance mismatches and crosstalk. For example, a 10GBASE-T PHY requires less power when operating with better performing connectors and cable. Reducing the noise and crosstalk in high-bandwidth, high-performance systems helps to

significantly reduce the latency, power and footprint, and cost of the silicon that must compensate for those impairments.

Conclusion

A host hardware offload solution can be partnered with cost-effective connectors that do not have to be as tolerant to thermal gradients. Where a system designer may have seen a significant increase in thermal properties and power when looking at transitioning from GbE to 10 GbE on an existing platform, an NTA solution makes the transition more manageable, possibly extending the life of existing systems and the use of existing components, like connectors.

Connector vendors can certainly make life easier for system designers who are researching and developing the next-generations of systems. Better impedance matching and reduced crosstalk can contribute significantly to enhanced system performance, reduced power requirements, reduced footprint and most importantly, reduced cost. Like Tehuti Networks host hardware offload method of making minor improvements with large gains, connector vendors can do the same.