

## Improving Web server performance

Due to the cost of network processing, high-end Web servers in the near future will be unable to handle the continually increasing demand of more clients ever hungrier for more content. In fact, it is already common place today to use multiple servers to host a single Website. Solutions such as multi-core systems or TCP Offload Engines have provided some relief, but are limited approaches because of the way they deal with memory.

However, a technique known as Remote Direct Memory Access (RDMA) has the capability to fully saturate high-speed networks while leaving the CPU free to do other tasks. RDMA not only moves the protocol processing to the network adapter, but also moves data directly from userspace. This eliminates the need for costly memory copies.

While RDMA has proved successful in the high performance computing realm, it is not used widely because of compatibility issues with existing network infrastructures. But a new adaptation of RDMA called iWARP enables RDMA over ordinary TCP/IP-based networks. (See the iWARP below for more details.)

Leveraging iWARP, the research team at the Ohio Supercomputer Center has created a module (`mod_rdma`) that enables the popular Apache Web server to use RDMA to send and receive client data.

"In our study, we have outfitted a Web server with 10 gigabit iWARP hardware and used client machines running iWARP software modified `wget` and Apache bench programs," said Dennis Dalessandro, networking researcher for OSC. "The result improves Web server performance, both in throughput and client request rate."

### Lead Researchers:

- Dennis Dalessandro  
Ohio Supercomputer Center
- Pete Wyckoff, Ph.D.  
Ohio Supercomputer Center

### Research Title:

RDMA Enabled Apache  
(`mod_rdma`)

### Funding Source:

- U.S. Department of Energy,  
ASC Program
- NetEffect, Inc.

### For more information:

[www.osc.edu/research/  
network\\_file/projects/rdma](http://www.osc.edu/research/network_file/projects/rdma)



## Downloading at warp speed

As network speeds have increased to 10 gigabits per second and beyond, today's CPUs have been unable to sustain the increased network processing requirements while at the same time meeting computational needs. However, specialty network solutions, such as InfiniBand, have long been available to solve this problem. The downside is the incompatibility with the existing TCP/IP based networking infrastructure that is common today.

iWARP, though, bridges the gap between high performance networking and TCP/IP compatibility. The term iWARP refers to a set of published protocol specifications that provide remote read- and write-access to user applications, without operating system intervention or intermediate data copies. The result is higher throughput and lower latency transfers. While hardware implementations of iWARP have begun to emerge, a software implementation is useful to serve as a transition mechanism and for protocol testing and research.

Another benefit provided by iWARP is single-side acceleration. In other words, only one end of a connection needs to have iWARP hardware to see local advantages, if the other side is equipped with iWARP. This is particularly attractive for the very common single-server, many-client scenarios.

Experiments conducted by Ohio Supercomputer Center researchers show that, with single-side acceleration, the sender system load drops from 35 percent to 5 percent, and receiver load drops from 90 percent to less than 5 percent, for 1 gigabit-per-second communication.

### Lead Researcher:

Pete Wyckoff, Ph.D.,  
Ohio Supercomputer Center

### OSC Project Team

- Dennis Dalessandro
- Ananth Devulapalli

### Research Title:

Software Implementation and  
Testing of the iWARP Protocols

### Funding Source:

Sandia National Labs

### For more information:

[www.osc.edu/research/  
network\\_file/projects/iwarp](http://www.osc.edu/research/network_file/projects/iwarp)