SysKonnect Application Breakthrough

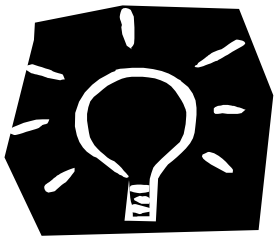

Making Gigabit-Ethernet Real: SysKonnect NICs Help Microsoft Break World "Land Speed Record"

Regardless of where system performance bottlenecks have occurred, historically, they have NOT taken place at the network interface card (NIC) level . . . at least not where **SysKonnect** products have been installed. Yet, how to break the elusive 1-Gig barrier?

Serving as a sort of digital Craig Breedlove, **Microsoft's** Jim Gray – senior researcher in Microsoft Research's Scaleable Servers Research Group and manager of Microsoft's Bay Area Research Center (BARC) – recently led top researchers around the U.S. to an "Internet2 Land Speed Record" using **SysKonnect** gigabit ethernet NICs for their test.

Using **Compaq** and **Dell** servers, and relying solely on *SK-NET GE-T* network interface cards from **SysKonnect**, the drive for full Gigabit throughput is now dramatically close to being available for everyday applications. The following notes, excerpted from Mr. Gray's findings, can be seen in their totality at http://research.microsoft.com/~gray/

In mid-February Y2000, teams from the **University of Washington** (Steve Corbato and others), **ISI-East** (Terry Gibbons and others), **QWest**, **Pacific Northwest Gigapop**, **DARPA's** *SuperNet*, and **Microsoft** (Ahmed Talat, Maher Saba, Stephen Dahl, Alesandro Forin, and team leader Jim Gray) collaborated to set a "land speed record" for TCP/IP. As such, these teams were the winners of the first ***Internet2 Land Speed Record***. The experiment connected two workstations with **SysKonnect** Gigabit-Ethernet network interface cards, via ten *SuperNet* hops (Arlington, NYC, San Francisco, Seattle, Redmond).

The systems delivered 750 mbps in a single stream tcp/ip (28 GB sent in 5 minutes) and about 900 Mbps when a second stream was used. This was over a distance of 5600 km, and so gives the metric 3 PetaBumps (peta bit meters per second). It was "standard" TCP/IP but had two settings: "jumbo" frames in the routers (4470 bytes rather than 1550 bytes) which gives the endpoints fewere interrupts, and also the window size was set to 20 MB (since the round trip time was 97 MS you need that much of a window to hold the "in flight" bits). The details are described in the submissions to the Internet2 committee, at http://research.microsoft.com/~gray/

ISI, **Microsoft**, **Qwest**, the **University of Washington**, **HSCC**, **PNWGP**, and **DARPA** demonstrated sending 8.4 GB worth of data in 81 seconds synchronously at a rate of 957 Mbps on two TCP/IP streams using Workstation class machines running *Windows 2000* on both ends. The source point is Redmond, Washington and the destination point is ISI-East at Arlington, Virginia via **DARPA's** *SuperNet* for a total distance of 5,626 Km. The experiments were done using the *Speedy* benchmark program running on *Windows 2000* at each endpoint. The program is a *WINSOCK* based application used for sending streamline data over a TCP connection. The experiments used the standard *Windows 2000* TCP/IP network stack running on **SysKonnect** Gigabit Ethernet cards.

- *more* -

SysKonnect Application Breakthrough

	Redmond-Arlington (MS-ISI)
User Data GB sent	8.4 GB
Elapsed time	81.359 seconds
User data rate Mbps	957.369 Mbps
Path taken	DC -> SEA (hop description) (hop #) (hop IP) ------- Arlington Virginia, ISI SysKonnect GbE 1 140.173.170.65 ------- Juniper M40 GbEArlington Virginia, ISI Interface ISI 2 205.171.40.61 ------- Cisco GSR OC48Arlington Virginia, Qwest DC Edge 3 205.171.24.85 ------- Cisco GSR OC48Arlington Virginia, Qwest DC Core 4 205.171.5.233 ------- Cisco GSR OC48New York, New York, Qwest NYC Core 5 205.171.5.115 ------- Cisco GSR OC48San Francisco, CA, Qwest SF Core 6 205.171.5.108 ------- Cisco GSR OC48Seattle, Washington, Qwest Sea Core 7 205.171.26.42 ------- Juniper M40 OC48 Seattle, Washington, Qwest Sea Edge 8 208.46.239.90 ------- Juniper M40 OC48Seattle, Washington, PNW Gigapop 9 198.48.91.30 ------- Cisco GSR OC48 Redmond Washington, Microsoft 10 131.107.151.194 ------- Redmond Washington, Microsoft SysKonnect GbE
Distance (KM)	5,626 km
Pteta bmps (bps*m)	5.38 Pbpms
Source& Sink application	speedy.exe
Source & Sink OS	Windows 2000
Source computer	Compaq SP750 dual 733 MHz, 256 MB RAM, **SysKonnect** GbE
Source computer price	~5k$
Sink computer	Dell 4400 dual 733 MHz, 1 GB RAM, **SysKonnect** GbE
Sink computer price	~6k$
Registry settings	Hkey_Local_Machine\system\Currentcontrolset\Services\Tcpip\Parameters TcpWindowSize = 20 Meg (20971520) Tcp1323Opts = 3 (Window Scaling)
MTU	4470 bytes

Table 1: **Setup and Configuration of Experiments**

The sender and receiver ran the *speedy* application documented in attachment 1 and 2, visible at http://research.microsoft.com/~gray/ using standard TCP/IP WinSock API's:

1) The sender formats 8.1 GB of data into a user specified buffer length (in KB) with unique random data, calculates the checksum in the user application, stores it in the outbound buffer, and sends the whole 8.1 GB worth of data on two TCP connections synchronously, and then terminates.
2) TCP does the flow control.
3) The receiver accepts the incoming data buffers, implements its own checksum in the user, over and above TCP's own Checksum built-in mechanism, on all incoming data for validating the content, and keeps track of how much data has been transferred and in how much time.
4) The elapsed time is measured as the time from the start of the receiver to the end of the receiver.
5) The receiver's console log for the experiment is in the PowerPoint presentation, again, at http://research.microsoft.com/~gray/

# # #

SysKonnect focuses on the development, production, and sale of high-end network products, including high-performance network interface cards (NICs) for Gigabit Ethernet, FDDI/CDDI and token ring, as well as FDDI concentrators for strategic networks. Headquartered in Ettlingen, Germany with locations in Great Britain and the USA, SysKonnect products are available via OEM partners and through the company's worldwide sales channels. Phone 408-437-3800 or visit www.syskonnect.com for more information.