GILDER TECHNOLOGY REPORT

Published jointly by Gilder Publishing, LLC and Forbes Magazine www.gildertech.com December 2001/Vol. VI No.12

Storewidth Peers

In this environment Equinix's Internet Business Exchanges are rapidly gaining market share against all comers

Like the Internet itself—and thousands of companies across the Telecosm—you have arrived at last, eyes blinking, wrists throbbing, down a carpal tunnel of digital bliss, looking for the incandescent light at the end. Enthralled with telecosmic visions, you search for seraphic web pages humming the sidereal symphony of a virtual New World. But you are hard to please, and something has gone awry. Where is that promised land? Where is that uplifting litter-free litry *Life After Television* where the customer is supreme? Where is that gildered age of glass and light, music and money? Why are **SBC**'s (SBC) copper cages valued at $129 billion and **Global Crossing**'s (GX) crystal cathedral of fiber under $1 billion? Why is the paradigm innovator **Avanex** (AVNX) at $8? Is this a simulated worst-case "mirror world" from the polymathic David Gelernter or a funhouse mirror from the poet of depression Jim Grant? Amid the 3.6 billion Googled pages of first choice glitz and klutz and clitz and hype and putz and glitter, the Web seems a no-man's land, as burkaed and borderless and anarchic as Afghanistan, with push-porn bats "swarming" up from the caves. You wonder where you are, and what you are waiting for.

Oh, well, in our deflationary concerns we sometimes wonder too. But if you are lucky, perhaps you will be coming to our Storewidth jubilee at the Ritz in Laguna Niguel, where liquidity seems limitless and golden in the sunlight by day and iridescently pink at dusk, and where Pacific breezes waft across the infoscape down broadband beaches, verdant fairways, and Ethernet boulevards. Basking in the warm sun and camaraderie, you can peer over the historic cliffs of costs toward the oceanic price elasticities of bandwidth and storage (all to be illumined there from March 24 to 27 at the second annual Storewidth conference).

Then, perhaps you think storing is boring and would rather sit at home or in your office, where you are merely perplexed by the paltry performance of your Internet connection, wondering why desired web pages do not pop up as swiftly as unwanted ones and why you cannot gain control of your information and your life. Your problem is that you—and the Internet—lack storewidth.

Defining the interplay between bandwidth and storage, I named the space last year, without working out all the specs. Consider it a beta. But look around and you can get an idea. You will see an immense global network of some 283 million miles of fiber-optic lines each with a potential capacity measurable in terabits (10 to the 12th)

In This Issue:

Cover Story: Post-diluvian data masses; Infinite bandwidth & storage converge; Parallel paradigm; Just listen (.com); Jitter bugs out; Scale Eight delivers globally; Google's speed secret; Equinix rising per second. That covers most of the bandwidth. The storage is just as fabulously voluminous (Figure 1). Scattered around the world you will see some billions of humble disk drives—JBODs by name (just a bunch of disks)—containing a total of tens of exabytes (10 to the 18th) of storage capacity.

All these realms of storage and bandwidth, iron and air, files and fiber, are rushing down a cliff of costs toward a nirvana of nearly nothing: a billionth of a cent for a room for a night in a bit motel (Figure 2); a few thousandths of a cent for a canonical bandwidth service called a DSO mile month.

Post-diluvian data masses

Since scant shares of the world's information—or your life's file—are wanted in real time, the essence of the Internet is the search, discovery, and recovery of stored data. Even a streaming video of the nightly news is stored on a series of buffers along the way. The heroic 5,500-fold rise in Internet traffic since 1994, documented by Larry Roberts among others, is chiefly a continuing expression of our ever-cheaper ability to store and share. Evidence from the high-end users surveyed by Keynote indicates that in the face of this traffic surge, average performance for those with a broadband connection has actually improved from a 12 second delay four years ago to a 3 second delay this year.

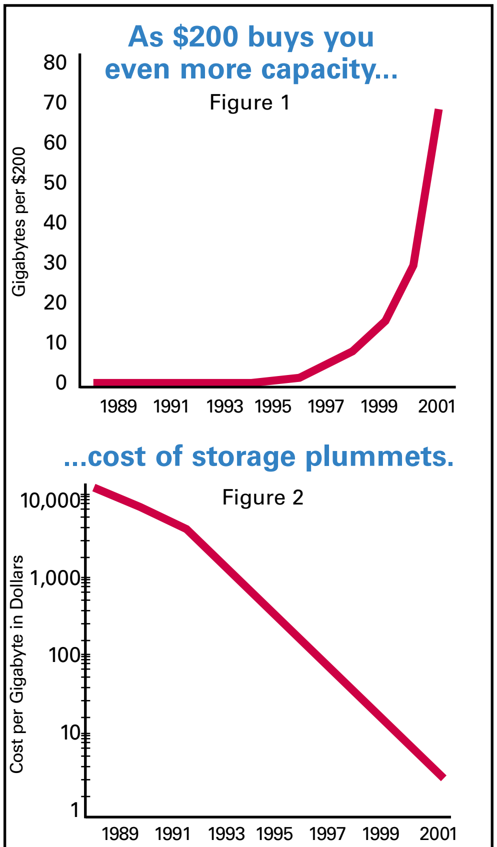

Source: WSJ and U.S. Department of Commerce

These reports signify a pace of advance in transport and switching technology that may be fairly described as stupendous. The *GTR*'s Charles Burger calculates that to sustain this year's Internet traffic with the communications technologies of six years ago would have cost $39 *trillion* or four full years of GDP. The bandwidth stars of the Telecosm have performed a job far more in line with the stock prices of the so-called bubble of the year 2000 than of the crash of 2001.

Today's Internet traffic continues on a pace of threefold annual growth. Dwarfing this Internet tide of moving bits, however, are the swelling troves of digital storage. In 1996, the bulk of storage was analog, inscribed on tapes, records, films, and microfiche that were sealed off from the Net. In that year, so guess the storage kings at **EMC** (EMC), digital storage facilities around the world held a total of 200 terabytes of data. Two hundred terabytes is a lot—forty kilobytes for every man, woman, and child on the face of the earth, including the caves of Afghanistan and the cafes of Great Barrington. By early 2001, however, converted to Internet-ready digital CDs, DVDs, disks, and tapes, total digital storage was 10 exabytes. A 50,000-fold increase in an industrial resource in half a decade must be unprecedented in the entire history of economics. And according to a Berkeley study, another 1 to 2 exabytes was added in 2001. That is the storage miracle.

I do not use exclamation points! But let me point with an exclamatory wave of my arms to these numbers. Distributed across the globe, ten exabytes translates into two billion bytes of digital information—some 2,000 big books of data—for every man, woman, and child on the face of the earth.

Numbers so large arouse skepticism. No one knows much about what the information is or why it was stored or where it resides, in what bloc, file, logical unit, disk array, track, rack, data center, warehouse, bitmine, cache, cave, cage, or other mnemonic device. Like most big books the data is seldom opened or read; 85 percent of it is said to be random, redundant, or obsolete. But hidden in its recesses and speleological mazes, in its caves and cages, disk farms and tape racks, are presumed to reside vast potential resources of knowledge for an information economy and, as we have lately come to recognize, elusive webs of clues to hidden threats of terror and paralysis.

What we can say for sure is that remote in your household and small business you cannot get swift access to these incomprehensible exafloods and data masses. You are probably happy that they cannot get access to you. Much of the pile is still opaque amid billions of these JBODs and it might as well be on the moon. Getting it off the moon and down that last long mile to you is one way to define the province of storewidth.

Infinite bandwidth & storage converge

At the Storewidth 2002 conference will gather many companies promising to address this challenge by accelerating the performance and curbing the latency of CDNs or content delivery networks and Web servers and storage switches and cable connections and digital subscriber lines. They will promise to Akamize your data in 13,000 caches or centralize it in four or strew it around the globe from satellites. **Narad Networks** will offer to pump it down new spectral overlays of gigabit Ethernet on current TV cables. In *Telecosm*, I wrote extensively about these ever expanding channels. But as the Renaissance visionary Gelernter of Yale and **Mirror Worlds** gently reminded me in a review and in a recent interview, much of the problem and the solution will come not in the last mile but in the household and business itself. The key to the response is addressing a new configuration of abundances and scarcities in which, for the first time in history, storage will be abundant.

The computer industry long ago learned to waste transistors to achieve faster performance and turn around time for its products. Slowly, but rather surely, the information industry is learning to waste bandwidth. However, as Gelernter observed, "Industry has yet to learn how to squander *storage* creatively."

Throughout most of the computer age, storage of all kinds has been oppressively scarce. The basements or attics of your computer held no extra room. The fabled Y2K crisis stemmed from a lack of memory for more than a two-digit date. Until the early 1990s, PCs had virtually no attic at all. Afflicting every **Microsoft** (MSFT) DOS PC was a memory ceiling of 640 kilobytes, reflecting Bill Gates's early view that no one could use more and conceding to **Apple** (AAPL) the markets above that barrier, chiefly its still vibrant franchise in graphics and pictorial applications.

The PC's basement of hard disk storage also was shallow and cramped. Until the early 1990s, hard drives advanced far more slowly than Moore's law, with areal densities increasing at a rate of only 27 percent per year. All archival storage had to go on tapes and much of it remained analog. Then came the explosion, with a set of synergies in platter quality and flatness, semiconductor thin-film fabrication, magneto-resistive drive heads, fast accurate servo motors, digital signal processing for bit detection, and related software all rushing forward in parallel to unleash a miracle in hard drives.

Doubling every 11 months—compared to a Moore's law cycle of 18 months—a high-end PC disk drive moved from 100 megabytes in the year 1992 to 120 gigabytes this year. In 1992, a 40-megabyte drive cost $150. In 2002, $151 gives you a 60 gigabyte GXP from **IBM** (IBM), which has already launched a 120 gigabyte model available to early adopters for $400. The 60 GXP represents a 1,500-fold drop in price per bit in a decade, while U.S. personal income nearly doubled.

Parallel paradigm

From the relative rates of progress on the three key vectors of industry advance—computer processing, bandwidth, and storage—emerges the outlines of a storewidth solution. Computer processing will advance at the pace of Moore's law. Optical bandwidth and storage will progress between 2 and 3 times faster. With microprocessors still operating on 32 bit data paths and migrating only slowly to 64 bits, with computer buses and peripheral connectors running at about a gigabit a second and stepping up to four gigabits in coming years, the bandwidth inside the personal computer is growing at a much slower rate than the bandwidth outside it on the increasingly optical network. In computer science, the answer to processor shortfalls is usually parallelization. So it will be with storewidth.

Rather than addressing data one step at a time, in a serial queue, as in the classic Von Neumann model, computer architects find ways to distribute the data across many processors and work on it simultaneously. In massively parallel computers, they optimize the processors less for fast execution of instructions at each processor than for broadband communications between processors. Even on the miles of microscopic wire on the backplanes of a single computer, the speed of light imposes its toll. The challenge is to maintain coherence (keep the data the same) between memory sites despite distance differentials between processors and shared storage and despite conflicts and collisions in a shared memory.

In computer science, the answer to processor shortfalls is usually parallelization. So it will be with storewidth.

Costly decades of federally subsidized supercomputer research in NUMA (non-uniform memory architecture) have led to a dismaying recognition of failure. No matter how many ingenious caching algorithms were deployed, a parallel multiprocessor is no more efficient than multiple computers operating in parallel, peer-to-peer, on a fast network. At first, a fast network was a mere megabits per second confined to a supercomputer center. Now with WDM optics, a thousandfold faster network has been spread out across the planet. The network is the computer and the computer is massively parallel, sharing storage around the globe. In such a NUMA configuration the key challenge is still speed of light latency and response time.

Optical bandwidth makes such a planetary system possible. But it cannot surpass the lightspeed limit. A rule of thumb in computer design ordains that bandwidth problems are solved in hardware, latency problems in software. To improve performance in latency and response is chiefly a problem of developing parallel software to tap the global scatter of abundant resources in storewidth. A 120 gigabyte drive can hold 60 full length MPEG 2 movies or 24,000 MP3 songs or 100,000 high resolution digital photos or scores of downloaded TV programs and home videos. But this abundance of content is valuable chiefly to the degree that users can share it. As home storage mounts, a peer-to-peer paradigm—computers in the home and business sharing resources across the network without a top down structure—is inevitable. But this model now collides with perplexing issues of copyright and intellectual property.

Just listen (.com)

Both in court and in the marketplace, the current strategies of the music industry are doomed to failure. In court, the music lawyers are trying to follow up their victory against Napster with a further campaign against the sharing of content across the Net. Yet Napster's free successors—now handling close to double Napster's peak 9.6 petabytes in February this year—are using **FastTrack**

PEERING INTO THE FUTURE

The Recording Industry Association of America (RIAA) blew it with its knee-jerk court victory over Napster. Customers fled Napster and dispersed across the Net. Over 3 billion files were downloaded using just the top four file-sharing systems in August compared to 2.79 billion files downloaded at Napster's peak last February. The most popular, FastTrack, chalked up 970 million downloads and 1.51 billion file transfers in September and 1.81 billion in October, an 87 percent increase in two months (**Chart 1**).

Emerging file-swapping services do not require the centralized downloadable file index that made Napster vulnerable to law enforcement. Peers can now independently distribute software and swap movies and games, in addition to MP3 files. Nevertheless the music and entertainment companies are heading back to court, nursed by pricey lawyers fighting against a human race stubbornly wanting to communicate. But to end post Napster Internet "piracy" the courts would have to shut down the Internet.

The power of human social interaction is often underestimated. In the early days of telephony, sociability was dismissed as idle gossip and discouraged until chit-chat began to fill company coffers. The Internet was launched in high seriousness by the DOD in the 1960s to connect computers. The first killer app: chit-chat by way of e-mail. The network connects people, not computers.

Humans prefer conversation to almost any other form of entertainment. As **Chart 2** shows, a measure of non-digital communications flows in 1999 yielded 15 petabytes of combined radio and TV content, while telephone and the postal service combined for a colossal 726 petabytes, most of it unique information. According to the U.S. Department of Commerce, in 1997 (the last year in which all relevant data was recorded) phone companies and post offices brought in $314 billion in revenues, while broadcast communications—radio, television, newspapers, and magazines—culled just $37 billion. Humans won't pay much for broadcast content. Almost all newspaper, TV, and radio revenues come from advertising of goods and services we *are* willing to pay for.

Cheap storage and abundant bandwidth dictate the architecture of the network. Decentralized peer-to-peer networks waste both bandwidth and cheap local storage capacity.

Streaming video startups such as Digital Entertainment Network, Pseudo.com, Pop.com, and Quokka Sports blame failure on a dearth of broadband last mile connectivity. Not so. These networks conserve storage and waste time, making little sense when cheap, abundant storage allows for any song, any movie, any game, any *time,* any *place*.

The value of interactive networks grows with scale. The value of a broadcast or point-to-multipoint network (e.g., newspaper, TV/radio, streaming content) is proportional to the number of users, while the value of a point-to-point communication network, according to Metcalfe's law, is proportional to the square of the number of users connected. Thus, a peered network with a million members may be on the order of a million times more valuable than a broadcast network with a million users (**Chart 3**).

-Charles Burger

Post-Napster file sharing surges past Napster's February peak.

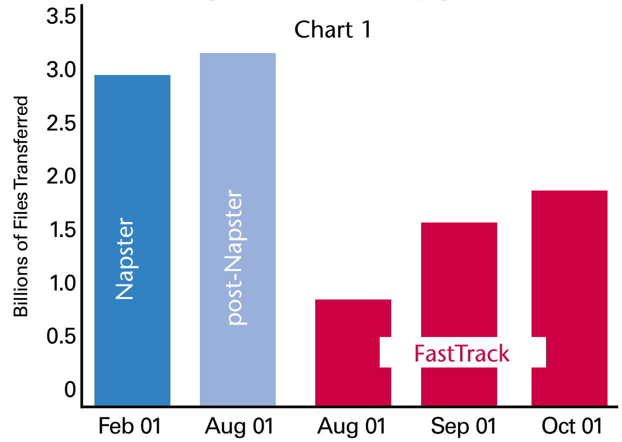

Source: Webnoize

Humans prefer communication to entertainment.

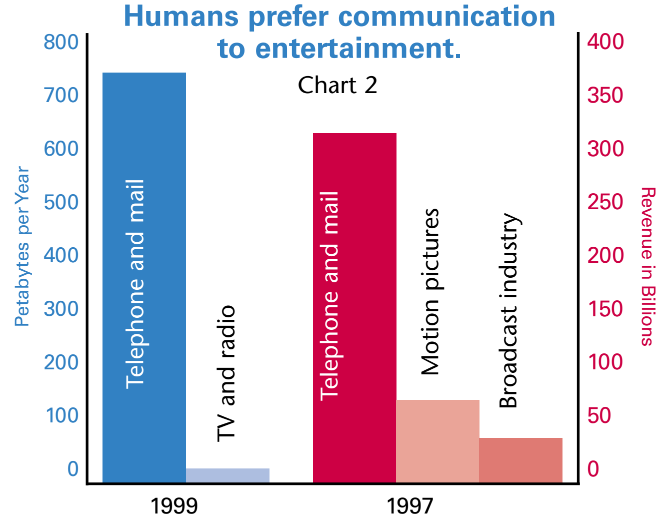

Source: University of California and Odlyzko, University of Minnesota

The value of peer-to-peer networks dwarfs broadcast networks.

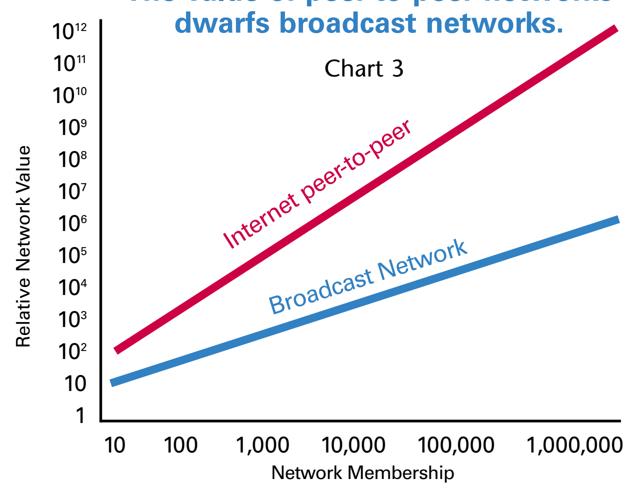

software from Holland that lacks the offending centralized index and the focus on MP3 files that incriminated Napster. Most of the new peer-to-peer file sharing systems are merely turbo-coded versions of ordinary file transfer software. To stop them the music industry must outlaw the basic structure of the Internet.

When your product is stolen by thieves, you have a police problem. When it is stolen by millions of honest customers, you have a marketing problem. At present, otherwise honest people download 150 million songs free over the Internet daily. That means 50 billion songs a year, the equivalent of three billion CDs of music or more than three times the 850 million copyrighted CDs sold annually. The "Napster" effect has roughly quadrupled the number of songs used in a year, expanded dramatically the range of music heard, and potentially enlarged the industry. The challenge now is to make it into a viable business.

The music industry's market responses to Napster—Press-Play, from **Sony** (SNE) and **Universal** and **Music.Net** from **EMI**, **Warner Brothers**, and **Bertelsmann**—are trying to compete with a free service with pay services offering less. They bar the sharing of music files, lists, and combinations that are the focus of current music users on the Net.

Offering a more promising scheme is Rob Reid, author in 1997 of the first definitive book on Internet commerce, *Architects of the Net*. Now aiming to become one himself, he has launched a company called **Listen.com**. It has created a system for distributing songs that both accommodates the current users and lays the foundations for a robustly profitable music business on the Net. Transmitting music files in unplayable form called residual audio, Listen.com allows the user to do whatever he likes with the file—storing, listing, filing, exchanging—until he actually plays the item. At that point he must qualify as a subscriber or a micro-payer in order to receive a real-time secure activation strip, containing the two percent of the file that renders the material playable. Designed for a world of storage and bandwidth abundance, where the prime impulse is sharing experiences, Reid's model makes his customers his marketers. It can spur the Net to new heights and create a new arena of profitable commerce.

Jitter bugs out

Just three years from now will begin the terabyte era, when a few hundred dollars will get you a terabyte drive, which could hold 480 theater resolution films, or 200,000 songs, and so on, in digital form. And that is in just one household or office desktop. Put them in 100 million households around the globe five years hence and you have 10 exabytes of storage just in people's homes. This is storage abundance, with nary a storewidth challenge in sight. Over the short hop between your disk and your display, these songs and images, films and games will play out utterly without latency or jitter.

Trying to achieve a similar effect out on the network, with less conspicuous success, is an array of contending companies. As Porter Stansberry pointed out at last year's Storewidth conference, these content delivery networks, "edgewidth" accelerators, and gigantic Web servers all aim to give users a low latency, jitter free experience. Yet, Napster-style users do not worry about latency or jitter. Their songs are on local storage. Nor do **TiVO** (TIVO) or Replay television customers, who download their favorite programs into what may now seem meager 40 gigabyte drives. Nor will the millions of users with hundreds of gigabytes of home storage. The entire industry strains paradoxically to deliver old, stored material in real time. But as Nicholas Negroponte predicted eight years ago, the bulk of Internet transmission will be store-and-forward. The medium for store and forward is peer-to-peer.

Scale Eight delivers globally

In a peer-to-peer architecture, data is served not through specialized storage servers, routers, and switches, whose optimized processors vainly strive to empty the bottleneck their very presence creates, but from your off-the-shelf PC, empowered by the optical Net and the terabyte era. Yes, commercial enterprises on a peer-to-peer Net will be supported by a recognizable infrastructure of servers, archival storage, data-mining facilities, and other technologies that will originate the product, protect its intellectual property, maintain archival libraries, and support and complement its use. Listen.com is deploying a server farm to deliver songs and strips to its subscribers. Beginning with 1.2 terabytes, the company plans to move to 5 terabytes shortly and use proxy servers at ISPs to cache popular songs. But the company has no current plans to use the storage facilities of EMC or **Network Appliance** (NTAP) or even **Scale Eight**. Instead, it is cobbling together its own cheaper JBODs on an **Oracle** (ORCL) 8i database and Linux servers.

When your product is stolen by thieves, it's a police problem. When it is stolen by millions of honest customers, it's a marketing problem.

Leading the way in parallel strategies for skirting the processing bottleneck is Scale Eight whose technology chief, 28 year old Joshua Coates, has concocted nearly one million lines of ingenious code to weave together a true system of globally redundant and accessible storage. It uses cheap disk drives in standard configurations rather than ever more expensive turbocharged RAID systems. Back to a humble but global JBOD. Obeying a fundamental law of the microcosm, that multiple, low-power, slower switching systems will outpace and outperform super-fast single points of frustration, the Scale Eight's system seems likely to transform the world of storewidth, from the enterprise to the global wide area network.

Coates disdains most of the complexities of storage area networks and storage optimized switches combining protocol, device, and application, and storage routers and

Fibre Channel esoterica that populate the acronym zoo. With hundreds of thousands of engineers pushing forward switches and routers and optical technology for Ethernet and TCP/IP, he is loath to rely on a few thousand engineers to advance the storewidth industry with proprietary storage switches and routers and other specialized devices. A "storage router" from **Cisco** (CSCO) costs about five times what an ordinary Cisco router costs and an Ethernet switch from **Extreme Networks** (EXTR) costs about one third of the price of a **Brocade** (BRCD) Fibre Channel switch. Meanwhile, on the storage side, IBM, **Seagate**, **Quantum** (DSS), **Intel** (INTC), **Texas Instruments** (TI), **Sun** (SUNW), and Microsoft are developing disk drives and controllers and other cheap devices. Always an advocate of parallelism, Coates wants these companies to specialize in their own fortes. He will develop software to integrate their products into a global storewidth network.

The most acute bottlenecks in the storewidth infrastructure are human beings. Coates wants to get them out of the way. Despite the advances in NAS and SANs from such companies as **Veritas** (VRTS), Network Appliance, **Hitachi** (HIT), and Brocade, the storewidth bottleneck will endure as long as humans remain responsible for pinpointing the location of data on the network.

RAID controllers and storage systems prevailed, in fact, by shielding humans from knowledge of the actual location (the disk, track, or bit cell) of the bits. Similarly transcending locality is virtualization through mirroring, striping, and metadata access. Similar to Gelernter's elegant Scopeware, developed to take humans out of the information retrieval equation, relieving them of their file-clerking tasks, peer-to-peer systems work precisely because no storage engineer need know where the data resides. In the radical equality of the most cultishly pure versions, the only location the humans need to know is "somewhere on the Net." Once you assert the need to know, the need for the data to be here in this hardware and not there in that lesser trove, the bottleneck begins to build.

Scale Eight's system seems likely to transform the world of storewidth, from the enterprise to the global wide area network

Nevertheless, storage remains an extremely manual process. Scale Eight is seeking to leverage the power of the network by breaking the file system out of the local box and putting it on a global file system that understands the network and makes the software, not the human, responsible for coordinating the troves of data.

Network Appliance and EMC took the first step in the right direction, enabling multiple storage appliances automatically to transfer copies of data from one device to another, or from one facility to another across a network. Hitachi, ascendant in the big storage appliance space, is smoothly integrating the SAN model with a network-attached system. Equiped with Sun software, Hitachi Lightening Storage Array is a shared memory device with 1.5 gigabytes of memory, 37 terabytes of storage, 512 disks, and multiple processors—essentially a parallel supercomputer that serves files.

Highly parallel data-flow movement through silicon is what enables **BlueArc** to zoom past the competition, serving up files at least 40 percent faster than its top competitors and enabling the sharing of a single machine by tens of thousands of simultaneously-connected users. Peered clusters of compute machines, working in parallel inside NAS units, power Vinod Khosla's storewidth sweetheart, **Zambeel**.

Google's speed secret

Moving outside of the box, parallelism based on peer-to-peer is everywhere changing the face of storewidth. Parallelism is what makes Google's search engine sort through 3.6 billion plus web pages and bang out the search results in typically less than one fifth of a second. Following the Coates principle of leveraging commodity devices, Google uses some 10,000 Linux PC servers clustered in parallel, each linked to 80 plus gigabyte drives. And a highly-parallel computing architecture is what is at play in Scale Eight's four storage centers (in CA, VA, U.K., and Japan), keeping their cheap off-the-shelf, commodity storage humming (without a RAID drive in sight). Several years ago, Dr. Andrew Grimshaw of **Avaki** showed that a parallel cluster of 256, 400 megahertz Pentiums with cheap IDE disks, running over a switched 100 Mbps Ethernet could serve up files nearly 40 times faster than a Network Appliance Filer of that era.

Most specialized storewidth box makers, accelerators, and engineers of Fibre Channel cul de sacs, are optimizing the scarcity—processing—rather than riding the abundances. The storewidth paradigm dictates wasting the increasingly free resources of bandwidth and storage to bypass the slow costive snarls of bits and wires, routers and servers, silicon and copper, switches and buffers, queues and protocols that currently populate the Internet. Overwhelming all tricky customized solutions, Ethernet moving to 10 gigabits per second, Infiniband unleashing fast servers on industry standard platforms, cheap storage multiplying its capacity into the terabytes per drive all will advance on their parallel paths. Moving beyond concepts of edge and center, the winners will be the companies that ride the wave rather than companies that try to modulate it for their own purposes.

Crucially, parallel systems transform the (relative) scarcity of processing power into an abundance of processors. Rather than trying to power through the server bottleneck with fewer, faster optimized storage processors, parallel systems exploit the abundance already in place: processors that are free because they are already there in the PCs of the world. Network attached systems are correct in exploiting the advance of the network. But they tend to conserve processors, linking scores of drives to single network appliances. Storage Area Networks offer vast amounts of storage in accessible forms, but they are restricted to a particular locality, with limited connectivity. Scale Eight's global reach and local action seems the model of the future.

Equinix rising

Meanwhile, investors should exploit the opportunity created by the current upheaval in the datacenter business. With the purchase of first **Digital Island** and then 30 **Exodus** data centers by **Cable and Wireless** (CWP) and AboveNet by **Metromedia Fiber** (MFNX), and with entry into the business by such telcos as **Genuity** (GENU), **Qwest** (Q), and **AT&T** (T), nearly all data centers are now owned by network players. Analogous to combinations of content with conduit, these hybrid centers seek to prevent content owners and networks from choosing their own optimal combinations of networks, peering points, and data centers. The data center managers must spend much of their time monitoring their network connections—and cross subsidies.

In this environment, **Equinix**'s (EQIX) Internet Business Exchanges (IBXes) are rapidly gaining market share against all comers. Focused on peering and exchange and on super secure facilities, Equinix has no networks at all, but invites leading providers to compete in connecting at its facilities. Among its customers and suppliers are Google which ascribes its fast parallel search capabilities in part to fast peering and redundancy at Equinix. Also aboard are Global Crossing, UUNET, **Yipes**, **Telseon**, **Williams Communications** (WCG), **Yahoo** (YHOO), IBM, Microsoft, and scores of financial companies attracted by security features modeled on the Federal Reserve. Developed by Jay Adelson, a co-founder of the Palo Alto Internet Exchange and thus a major figure in the evolution of Internet peering, Equinix was designed as a neutral hosting center that can accommodate diverse networks. Equinix is the only hosting company that posted a sequential revenue gain in the last quarter and it is rapidly becoming a crucial pivot of Internet infrastructure.

The government sees Equinix as a threat to monopolize the Net. Yet the market sees it as a financial convalescent and values its ingenious hubs at a total of $109 million. The market is crazy. Equinix will have $80 million of unrestricted cash on its balance sheet at the end of the year and it recently renegotiated its existing $150 million credit facility and reaffirmed its fully funded plan to become EBIDTA positive in the fourth quarter. A presenter at both Telecosm and Storewidth 2001, Equinix joins our list this month along with the promethean private company Narad Networks profiled in the June 2001 *GTR*.

George Gilder with Mary Collins Gorski
December 13, 2001

A digital fountain

Possibly inflicting a final blow to the "edgewidth" model of costly Internet bypass networks and accelerators is a new method of sending flawless files through noisy channels. Based on a new "holographic" algorithm invented at Berkeley by mathematician Mike Luby, a company called **Digital Fountain** offers a system for downloading any file perfectly regardless of interruptions, sequence, jitter, or lost packets, with just 5 percent bandwidth overhead. The secret is yet another form of parallelism. All packets are created equal, and as long as the recipient gets a full complement plus 5 percent, the file is flawlessly transferred. Rather than sending the content itself, Digital Fountain puts it into a mathematical blender that homogenizes the packets. While the laws of entropy prohibit reseparating blended eggs, the laws of mathematics–and Luby's algorithm–allow the receiver to perfectly reconstitute the original file from the transmitted metadata. Launched with support from Garage Technology Ventures and yours truly and backed by Sony and Cisco and an array of top line venture capitalists, Digital Fountain is transforming the parameters of file transfer on the Net.

Consider the storewidth dilemma faced by **Siebel Systems Inc.** (SEVL) developing customer response management software in San Mateo and Ireland and shipping the "builds" across the globe every night. Beginning with a T-1 line, the transfer took seven hours at a pace of 500 kbps. A seven-hour time change made the process intolerable. Hey, this is the age of bandwidth, the Siebel people figured, "throw bandwidth at the problem." So over a period of weeks, they acquired a 30-megabit per second (fractional E-3) line to Ireland and sent off the nightly build. Huh? The transmission still traveled at 500 Kbps and took seven hours. Ignoring speed of light latency, which consumed 200 milliseconds for every round-trip TCP acknowledgement and required to-and-fro retransmission for every lost packet, Siebel was trying to solve a storewidth problem with bandwidth.

Siebel discovered that Digital Fountain servers together with less than one megabyte of client software could remove the latency. Gone were all the delays and back channel acks and reacts. The result was a fivefold acceleration of the transfer. With all packets equal, lost packets do not entail retransmission, merely a proportionate extension of the flow. If faster transfer is needed, multiple streams can be sent in parallel with no change in the file and no need to get them in order at the other end. If you want to do streaming audio or video, you break up the file into segments and send each of them to be buffered at the other end. Recipients can join the stream at any time. If you want to filter a file through the troposphere from a geosynchronous satellite or bounce it around the globe through 25 routers that routinely relieve congestion by dropping packets, you need no back channel signals or retransmissions to assure a perfect transfer.

Over the next five years or so, the Internet core will turn into an all-optical fibersphere which will convey data on massively parallel wavelengths without dropped packets and multiple hops, jitter and jumble. However, wireless, satellite, and other noisy last mile channels assure that even then, the Digital Fountain technology affords a cheap and effective way to get the most out of the storewidth jungle.

- GG

TELECOSM TECHNOLOGIES

ASCENDANT TECHNOLOGY	COMPANY (SYMBOL)	REFERENCE DATE	/ PRICE	NOV '01: MONTH END	52 WEEK RANGE	MARKET CAP
FIBER OPTICS						
Optical Fiber, Photonic Components	Corning (GLW)	5/1/98	13.64	9.43	6.92 - 79.75	8.9B
Wave Division Multiplexing (WDM) Components	JDS Uniphase (JDSU)	6/27/97	3.63	10.08	5.12 - 76.63	13.4B
Adaptive Photonic Processors	Avanex (AVNX)	3/31/00	151.75	7.26	2.70 - 97.50	482.4M
All-Optical Cross-Connects, Test Equipment	Agilent (A)	4/28/00	88.63	27.27	18.00 - 68.00	12.6B
Tunable Sources and WDM Components	New Focus (NUFO)	11/30/00	20.31	4.04	2.10 - 62.88	305.0M
Crystal-Based WDM and Optical Switching	Chorum (private)	12/29/00	—	—	—	—
WDM Metro Systems	ONI (ONIS)	12/29/00	39.56	7.12	3.50 - 69.75	994.5M
WDM Systems, Raman	Corvis (CORV)	3/30/01	7.03	3.06	1.19 - 45.13	1.1B
Metro Semiconductor Optical Amplifiers	Genoa (private)	3/30/01	—	—	—	—
Optical Processors	Essex (ESEX.OB)	7/31/01	5.90	6.10	1.88 - 7.45	31.0M
LAST MILE						
Cable Modem Chipsets, Broadband ICs	Broadcom (BRCM)	4/17/98	6.00*	43.99	18.40 - 148.50	11.5B
S-CDMA Cable Modems	Terayon (TERN)	12/3/98	15.81	12.06	2.36 - 18.88	827.7M
Linear Power Amplifiers, Broadband Modems	Conexant (CNXT)	3/31/99	13.84	14.89	6.57 - 33.00	3.8B
Broadband Wireless Access, Network Software	Soma Networks (private)	2/28/01	—	—	—	—
Gigabit Ethernet Coaxial Cable Networks	Narad Networks (private)	11/30/01	—	—	—	—
WIRELESS						
Satellite Technology	Loral (LOR)	7/30/99	18.88	2.30	1.03 - 6.34	770.5M
Low Earth Orbit Satellite (LEOS) Wireless Transmission	Globalstar (GSTRF)	8/29/96	11.88	0.65	0.20 - 2.38	71.8M
Code Division Multiple Access (CDMA) Chips, Phones	Qualcomm (QCOM)	7/19/96	4.75	58.72	38.31 - 107.81	44.9B
Nationwide CDMA Wireless Network	Sprint (PCS)	12/3/98	7.19 *	24.95	15.72 - 33.25	24.6B
CDMA Handsets and Broadband Innovation	Motorola (MOT)	2/29/00	56.83	16.54	10.50 - 25.13	36.8B
Wireless System Construction and Management	Wireless Facilities (WFII)	7/31/00	63.63	5.19	3.31 - 52.11	243.7M
GLOBAL NETWORK						
Metropolitan Fiber Optic Networks	Metromedia (MFNX)	9/30/99	12.25	0.56	0.25 - 19.94	432.1M
Global Submarine Fiber Optic Network	Global Crossing (GX)	10/30/98	14.81	1.27	0.38 - 25.88	1.1B
Regional Broadband Fiber Optic Network	NEON (NOPT)	6/30/99	15.06	4.30	2.16 - 19.94	91.7M
National Lambda Circuit Sales	Broadwing (BRW)	6/29/01	24.45	9.52	7.50 - 28.88	2.1B
Internet Backbone and Broadband Wireless Access	WorldCom (WCOM)	8/29/97	19.95	14.54	11.50 - 23.50	43.0B
STOREWIDTH						
Java Programming Language, Internet Servers	Sun Microsystems (SUNW)	8/13/96	6.88	14.24	7.52 - 48.13	46.2B
Network Storage and Caching Solutions	Mirror Image (XLA)	1/31/00	29 .00	2.60	1.00 - 14.13	295.0M
Remote Storewidth Services	StorageNetworks (STOR)	5/31/00	27.00*	6.87	3.65 - 55.25	668.3M
Hardware-centric Networked Storage	BlueArc (private)	1/31/01	—	—	—	—
Virtual Private Networks, Encrypted Internet File Sharing	Mangosoft (MNGX.OB)	1/31/01	1.00	0.82	0.34 - 3.00	22.1M
Massively Parallel Global Storewidth Solutions	Scale Eight (private)	8/31/01	—	—	—	—
Secure Internet Business Exchanges	Equinix (EQIX)	11/30/01	1.65	1.65	0.33 - 7.63	132.1M
MICROCOSM						
Analog, Digital, and Mixed Signal Processors	Analog Devices (ADI)	7/31/97	11.19	42.50	29.00 - 64.00	15.4B
Silicon Germanium (SiGe) Based Photonic Devices	Applied Micro Circuits (AMCC)	7/31/98	5.67	13.63	6.01 - 88.25	4.1B
Programming Logic, SiGe, Single-Chip Systems	Atmel (ATML)	4/3/98	4.42	8.25	5.48 - 18.44	3.8B
Single-Chip ASIC Systems, CDMA Chip Sets	LSI Logic (LSI)	7/31/97	15.75	16.25	9.78 - 26.10	6.0B
Single-Chip Systems, Silicon Germanium (SiGe) Chips	National Semiconductor (NSM)	7/31/97	31.50	30.13	17.13 - 35.10	5.3B
Analog, Digital, and Mixed Signal Processors, Micromirrors	Texas Instruments (TXN)	11/7/96	5.94	32.05	20.10 - 54.69	55.5B
Field Programmable Gate Arrays (FPGAs)	Xilinx (XLNX)	10/25/96	8.22	36.11	19.52 - 59.25	12.1B
Seven Layer Network Processors	EZchip (LNOP)	8/31/00	16.75	6.23	2.70 - 21.19	45.4M
Network Chips and Lightwave MEMS	Cypress Semiconductor (CY)	9/29/00	41.56	23.02	13.72 - 29.25	2.8B
Field Programmable Gate Arrays (FPGAs)	Altera (ALTR)	1/31/01	30.25	22.76	14.66 - 34.69	8.8B

ADDED TO THE LIST: EQUINIX AND NARAD NETWORKS

* INITIAL PUBLIC OFFERING

NOTE: The Telecosm Table is not a model portfolio. It is a list of technologies in the Gilder Paradigm and of companies that lead in their application. Companies appear on this list only for their technology leadership, without consideration of their current share price or the appropriate timing of an investment decision. The presence of a company on the list is not a recommendation to buy shares at the current price. Reference Price is the company's closing share price on the Reference Date, the day the company was added to the table, typically the last trading day of the month prior to publication. Mr. Gilder and other GTR staff may hold positions in some or all of the stocks listed.

Gilder Technology Report ***Published by Gilder Publishing, LLC and Forbes Inc.***

291A Main Street • Great Barrington, MA 01230
Tel: (888)484-2727 • Fax: (413)644-2123 • Email: info@gildertech.com

EDITOR: George Gilder
PUBLISHER: Richard Vigilante
ANALYSTS: Charles Burger, Mary Collins Gorski, Bret Swanson
RESEARCH: John Hammill
MANAGING EDITOR: Debi Kennedy
DESIGNER: Julie Ward

SUBSCRIPTION DIRECTOR: Rosaline Fernandes
GENERAL MANAGER, GILDER TECH.COM: David S. Dortman
PRESIDENT: Mark T. Ziebarth