

GILDER TECHNOLOGY REPORT

Published jointly by Gilder Publishing and Forbes Magazine

www.gildertech.com

June 2000/Vol. V No. 6

The End of Storage

In storewidth, undershoot is still the problem. The industry is nowhere near supplying fast and seamless access to diverse Web resources

Bounding and Boeing across the silicon continent, from Akamai (AKAM) to Exodus (EXDS), from Novell (NOVL) to Network Appliance (NTAP), with side trips to the phenomenal Avanex (AVNX), we found ourselves nostalgic for the old concept of storage. Where can we find a capacious port amid the relentless storms of paradigm change? But no such luck. In all these stations on the paradigm path, none of the estimable executives and engineers—from Akamai guru Avi Freidman to the encyclopedic Exodus CEO Ellen Hancock—had anything much to say about storage.

Storage is already free. By the signal of plummeting costs for a key factor of production, storage is already over, priced at a billionth of a cent or less for a motel bitroom for a night. More than at a global convention of Christian evangelicals, university presidents, or welfare rights protesters, the insistent buzz is not for storing (hoarding is the word for that) but sharing.

Share this. In a world of pullulating polyglot data, the pressing agenda is not archival. It is the emerging new architecture of the Internet. It is storewidth—the conversion of abundant bandwidth and heterogeneous petabytes into accessible information. A rainbow coalition of digital diversity, the storage petabytes show up in varying forms, from simple text files to MP3 music to HDTV images to relational database queries to transactional interactions. These hide behind a magpie's nest of operating systems comprising a NUMA (non uniform memory architecture) system, which means they vary drastically in distance from the user and in bandwidth of connections, and thus they pay varying fees at the speed of light tollgate. In this maze, actual storage is the simplest part. Storewidth—the accessibility challenge—is at the heart of the next phase of Internet development.

Storewidth remains grossly inadequate. The average Internet packet makes seventeen hops between routers before arriving at your browser. Web pages contain as many as twenty-five trinkets or objects to be moved through the seventeen-hop gantlet. As a result, streaming video still comes in jittery gushes. Voice over IP is fitful and distorted. Access to Web pages evokes the world wide wait. The ideal of instant broadband images and transactions is still far from fulfillment.

In This Issue:

Cover Story: Hollowing out the computer; Cowboy storage; EMC's tetherless disk arrays; Procom sweeps storage; Mango tango

CDMA Wins in North America Page 4

Procom and Storage Networks
Join the List Page 8

Scores of thousands of firms are pursuing this trillion dollar challenge with frenzied ingenuity and creativity. Contrary to analysts who see the new technology as a bubble, our experience visiting companies and conferences shows us that the new economy is far richer and deeper and more efflorescent with clustering troves of enterprise than I have readily imagined. After centuries of misanthropic economists treating humans as mass men or assembly line cogs or gaping mouths, the power of individual minds is now detonating across the Net.

But all too often the technology stars hide behind veils of secrecy while maneuvering for vantage in the mazes of the IPO process, with its bizarre quiet period, its lascivious roadshow tease, and its elastic window seductively opening and then snapping shut on the fingers of the unwary or avaricious.

One of the most impressive companies in the storewidth space, for example, **Storage Networks** of Waltham, Massachusetts, plans an IPO on June 27 and faces the omerta of a quiet period. Meanwhile, it is making like so many New England clams. But they could not stop us from making clam chowder at Supercomm with Storage Network's sometime supplier ONI Systems. ONI volunteered that Storage Networks is providing one of the first markets for paradigmatic lambda circuits—wavelengths devoted to a single link from a customer to a storage center on the Net, or between storage centers around the globe. Moreover, clams (Companies Laboring Against Mandated Silence) cannot talk back, so we can say whatever we want about them, beginning with changing their name. Storage Networks is cumbersome so we will dub it SNI. Moreover, I cleverly acquired its S-1 and learned some key company secrets which I will divulge here. For example, the S-1 reveals that “our quarterly revenues and operating results may fluctuate in future periods...” They also confide that “there is no guarantee that Year 2000 issues will not in the future have a material adverse effect on our business...” Don't ever say we didn't warn you.

The SNI strategy feeds on the voltage of perhaps the most powerful differential in technology, the growing gap between bandwidth outside the computer and bandwidth within it. For most of the digital age, the speed gap between fast computer input-output (I/O) channels and sluggish network connections condemned the computer to live in a box. Despite the greater convenience of dispatching various functions—printing, storage, the applications themselves—to optimal points on the network where they could serve many people at once, everything was packed into a motherboard and a pastel chassis. Forcing this Mother Hub-board architecture were specialized, very short distance I/O channels which provided the only way to get to the CPU on time. Typical were FDDI (fiber distributed data interconnect or “fiddy”), SCSI (small computer systems interconnect or “skuzzy”), and their successors. Running together at up to 100 megabits per second, FDDI and SCSI excelled the shared bandwidth of network Ethernet by as much as one hundred to one.

Thus in those evil days, almost all storage was tethered directly to a particular computer and accessible only through it. While the network was driven by the power and possibilities of sharing, the data remained frustratingly distant from the network as multiple clients queued up for the attention of a CPU preoccupied with more urgent matters than spinning its disk drives in search of a missing byte.

Yet even as early as 1992, the shared storage market was already \$1 billion. As Dave Hitz, Vice President of Network Appliance, which invented Network Attached Storage, pointed out, “When was the last time you heard of someone paying \$1 billion to go one hundred times slower?” Bob Metcalfe's law of networks was the reason (see his new book out this month from IDG). The entire Internet was built on Metcalfe's Law. You are willing to go one hundred times slower in order to access a million times more data.

Now comes the revolution. The standardization of Gigabit Ethernet today gives the network rough parity with Fibre Channel I/O. By the year 2002, the standardization of 10 Gigabit Ethernet, impelled by **Broadcom's** (BRCM) one chip 10 gig transceivers, will totally reverse the relationship between internal and external bandwidth. Beginning even this year, companies will pay a 5X penalty in access speed for the privilege of keeping a function in the computer box rather than committing it to the network.

Hollowing out of the computer

In a 1995 email, Eric Schmidt, now CEO of Novell (NOVL), a company devoted entirely to storewidth, dubbed this effect the “hollowing out of the computer.” When the network is as fast as the computer's internal links, the machine disaggregates across the Net into a set of special purpose appliances. First to move was the printer, which became a fully networked device courtesy of Netware and **HP** (HWP). The next network appliances—displays and keyboards—moved onto the net in the form of X-terminals and **IBM** (IBM) or **Televideo** (TELV) workstations. But these precursor appliances normally used only enough bandwidth to transfer keystrokes and register ASCII code on a screen.

The real hollowing commenced with the move of storage to the Net. Comprising some ninety-eight percent of the transistors and other electronic domains in the system, storage in its many forms of buffers and registers and caches pulses at the very heart of computing. A computer may be plausibly described as an I/O device for various storage systems. You move storage onto the Net, and the Schmidt computer—as hollow as a CRT—becomes the prevailing architecture. The network becomes the computer at last.

Taking this trend and blasting it all the way out of the enterprise onto the Internet itself, Storage Networks has launched a software array that duplicates on the Net all the necessary features of local storage. It names the programs “PACS” because they offer Protection, Availability,

Continuity, Scalability, and Security of data. These benefits come through a front-end Virtual Storage Portal, based on Java, for monitoring and provisioning storage on demand, and through a back-end operating system for managing twenty-four hour command and control of a hierarchical disk and tape cascade. The company's one hundred fifty storage engineers (presumably more by now) can create storage equally accessible from your office or a data center from AT&T (T), Exodus, or **Global Crossing** (GBLX). Since one of its customers is X:Drive of Los Angeles, which rents disk space to individuals, SNI is indirectly serving residences as well. In the belief that keeping vital data within the walls of your company will soon seem as quaint as keeping your money under a mattress, SNI is on target with the paradigm.

If you are tying together a worldwide web of storage facilities bearing the very lifeblood of your business, you want to be up all the time; you want dedicated circuits, not a statistically muxed bitstream. Storage Networks, thus, also partakes of the paradigmatic mandate of wasting bandwidth with lambda circuits. It is not a bandwidth company (it assumes the continued explosion of bandwidth) and it is not a storage company (it assumes the rapid advance of storage capacities). It is a storewidth company, combining bandwidth and software that render storage as accessible and robust, secure and available, as a good bank. Its investors include Global Crossing, Exodus, Network Appliance, Veritas (VRTS), Hewlett Packard, Dell (DELL), and a host of venture capitalists.

Cowboy storage

Storage Networks began in 1998 by supplying storewidth for Vastar (VRI), an oil company in Houston owned eighty-four percent by Amoco. Four times a year Vastar had to make bids on oil leases based on terabytes of seismic data on tapes from oil fields around the world. Carting the tapes around the city to various company facilities in pickup trucks, the Vastar people began to think there might be a better way.

"Put all the storage in Houston," they told Storage Networks founder and CTO William Miller, "and link it to Calgary and Alaska and Saudi Arabia. Make it instantly available while we bid on the leases, and you're on."

SNI got its Houston network from an old codger and oilman named Lee Cook. In premature partnership with Bing Crosby in the 1950s to bring video on demand, Cook had acquired all the rights-of-way for a network around Houston. Miller and Cook made a deal, and the business was under way. Seismic data has much in common with rich streaming Internet data. Both are bulky and serial, and clients for seismic data require fast access. Meeting the challenge of the oil companies prepared SNI to fulfill the demands of complex web portals dispensing a variety of demanding datatypes.

Seeing the business take off in Houston under Miller, CEO Peter Bell, formerly of EMC (EMC), set up

in New York to target customers in the financial community, swamped in mires of mandated records, transactional demands, and data mining opportunities. Bell launched a storage center at AT&T's hub at 111 Eighth Avenue, connected by a fiber ring to the World Financial Center and linked to New Jersey under the Hudson River through 864-fiber cables from **Metromedia Fiber** (MFNX). The cables ran through the Holland and Lincoln tunnels to Internet nodes at Weehawken and Jersey City, which in turn were linked to the SNI control center in Waltham and to the SNI center in London via Global Crossing lines. The word is that **Merrill Lynch** (MER) now has some 14 terabytes of data with the company.

The drastic change in the networking environment gives Procom a huge opportunity to sweep into the 600 petabyte arena for storage...

Oil and finance are applications from the old economy. But also showing up in Waltham was **Yahoo** (YHOO) CTO Farzad Nazem who was faced with integrating such disparate companies as **Geocities** and **Broadcast.com** into its system. Next door in Waltham, **Lycos** (LCOS) faced similar challenges of combining diverse new acquisitions that demanded fast response times and robust transactions with jitter-free video streams and data mining of click streams, all totaling many terabytes of storage. Today most Storage Network customers are dotcoms.

The key issue for the company is barriers to entry. What SNI pioneered in two years can presumably be accomplished more quickly by other companies that benefit from the SNI learning curve. However, it will not be easy. Raising \$203 million of venture capital, hiring several hundred storage engineers, contriving two complex software systems suitable for diverse customers in oil, finance, and on the Web, negotiating thirty-six-month guaranteed Service Level Agreements (SLAs) with scores of customers from Merrill Lynch to X:Drive—all in just two years—Storage Networks is showing real prowess running with the bulls on Internet time. Although the field is crowded and there is scant room for mistakes, Storage Networks is the prime storewidth service provider today.

To achieve the PACS goals, SNI integrates gear and supporting services from such companies as EMC and **Sun** (SUNW) to Network Appliance and **LuxN**, and uses bandwidth from Telecom stars Global Crossing, Metromedia Fiber, and **Yipes**. To explain SNI's strategy—and the current state of storewidth—SNI founders turn to examples from the history of EMC down the Massachusetts Turnpike in Hopkinton. EMC remains the Goliath of the storage field.

CDMA WINS IN NORTH AMERICA

Though half the mobile phone subscriber base in North America is still analog, the future, and all the growth, is in the contest for digital supremacy. In 1999 CDMA began to pull away.

At the end of 1997, effectively CDMA's first year in the North American market, TDMA, with a two-year head start, accounted for almost 50 percent of all digital subscribers. And analog still accounted for almost 90 percent of the 54 million total mobile subscribers.

In the next two years CDMA surged, growing some 253 percent from Q4 '97 to Q4 '98, and another 154 percent through Q4 '99, almost double TDMA's 1999 increase of 88 percent. Nearly every major CDMA carrier experienced triple digit digital growth last year; no large TDMA carrier did so. By December 1999, CDMA had just edged out TDMA as the leading digital technology (Charts 1 & 2).

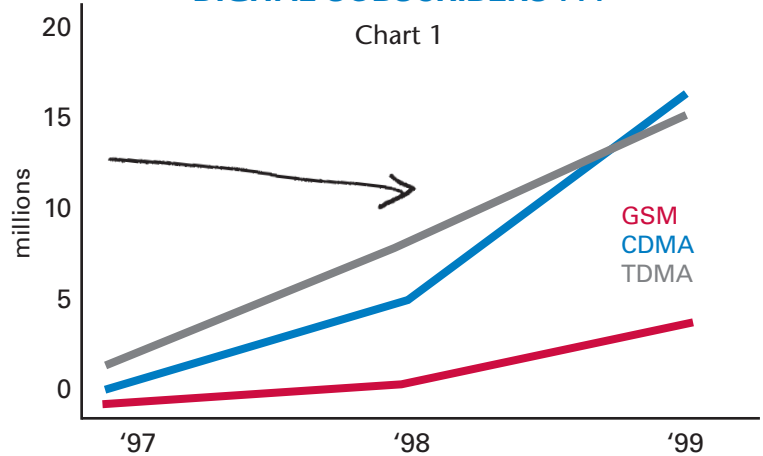
With half the total market still analog, CDMA's share of the total market was still only 20 percent at year-end. But handset sales, a leading indicator of subscriber share, suggest that as analog disappears and digital sweeps the field, CDMA will pull away from the pack. In 1998 TDMA handset sales still led CDMA by 1.4 million units. But in 1999, CDMA handset sales took the lead over TDMA by 2.4 million units, for a swing of almost 4 million units and 10 points of digital market share in 12 months (Charts 3 & 4).

Buoyed by its analog remnant, AT&T's share of all subscribers has remained fairly steady, hovering around 11 percent. But in the digital market, where all the growth is, AT&T's collapse has been as catastrophic as Sprint's rise has been meteoric. Unofficial Q1 numbers show the trend continuing, with Sprint adding an industry leading 876,000 net new subscribers, and AT&T adding about half that, for a total of roughly 10 million (Charts 5 & 6).

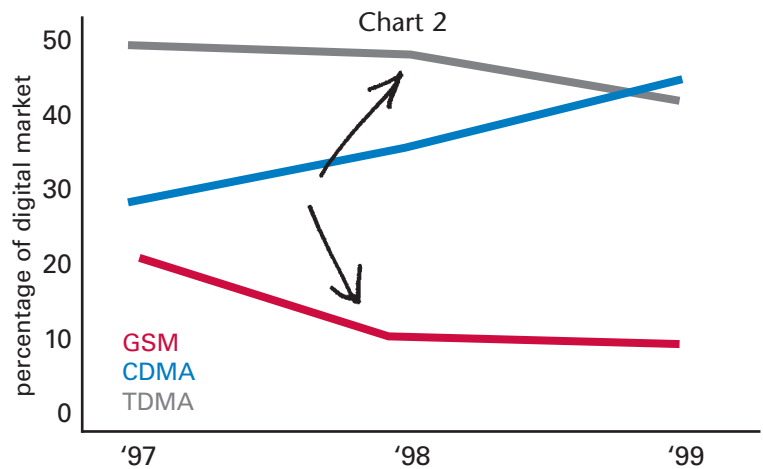
AT&T still easily outpaces all its major rivals in average revenue per subscriber, hovering around \$66 per month, \$12 to \$15 more than Sprint or Bell Atlantic. But in a market now driven by furious price competition, and facing greater than expected demand for data services, that may not be good news for a carrier whose capacity-strained technology makes it less able to benefit from elasticities of demand. **GTR**

Sources: Industry Sources, Dataquest, CDMA Development Group, EMC World Cellular Database

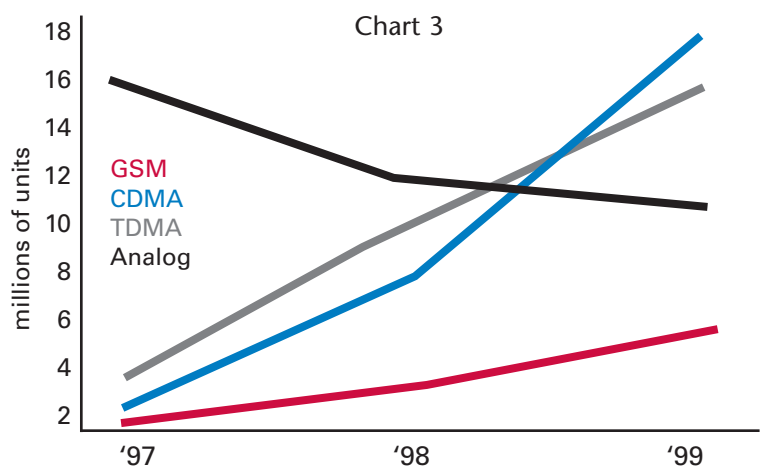
CDMA EDGES TDMA FOR LEAD IN DIGITAL SUBSCRIBERS . . .



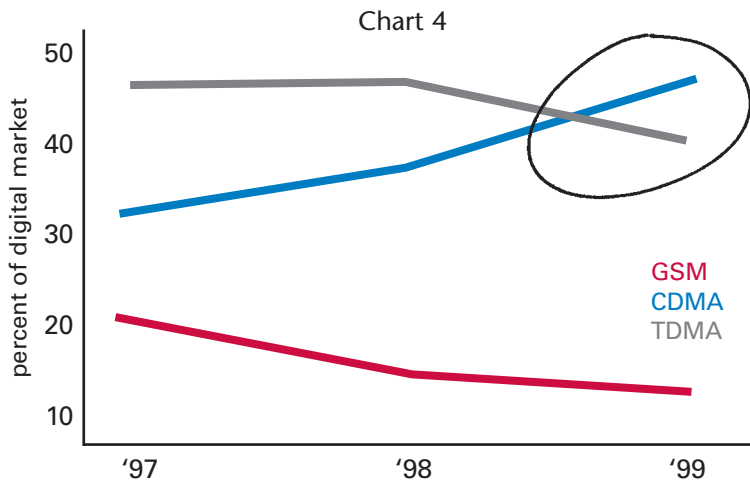
. . . AS TDMA, GSM MARKET SHARE PLUNGE



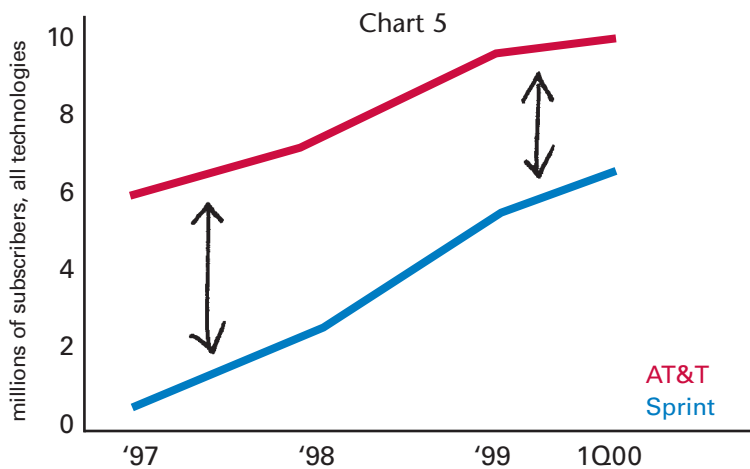
CDMA HANDSET SALES PASS ANALOG, TDMA



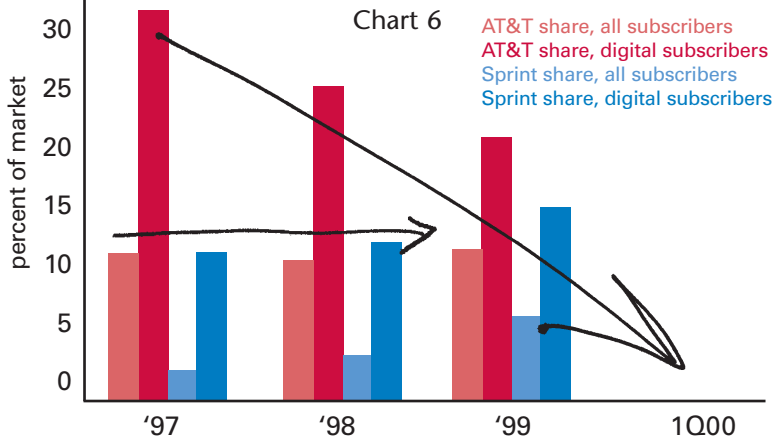
DIGITAL HANDSET SHARE PREDICTS GROWING CDMA LEAD IN SUBSCRIBERS



CDMA-ONLY SPRINT IS FASTEST GROWING CARRIER EVER



AT&T'S SHARE OF BOOMING DIGITAL MARKET COLLAPSES



EMC's tetherless disk arrays

In the mid-1980s, EMC began with the problem of multiple disk drives. Usually tethered to one processor in a single computer, with a multi-purpose operating system, disk drives were all-too-typical children of harried and hapless single parents. Raised by part-time operating systems, the drives seemed never to come first: They grew up sub-optimized, their data strewn hither and yon, unable to set priorities, make schedules, show up on time, or grasp the expectations of the working world.

EMC declared that it takes a village. It began as a Vista program for disadvantaged drives, taking cheap devices from Seagate and the Japanese, teaching them teamwork, and getting them jobs in plug-compatible storage systems for IBM mainframes. Within five years, EMC took sixty percent of the IBM storage market. Then in 1996, it launched a series of storage area networking (SAN) products that liberate storage arrays from particular servers and computer systems. A specialized back channel network, a SAN links disks and tapes, memory caches, load balancers, file servers and fail-over redundancy schemes and manages them as a unit, which in turn can be connected to the Web.

The SAN village seems the antithesis of the paradigm, a mortal violation of the Telecom's primary directive: waste bandwidth to conserve processing. SANs, in this scenario, conserve network bandwidth by spending processors like an Intel wet dream, contriving network detours tangled with smart hubs and smarter switches, specialized cables, extenders, and adapters far more complicated and expensive to build or manage than the LAN itself.

By contrast, NAS (Network Attached Storage) appliances attach directly to the Net, use mostly the same cheap and easy Ethernet connections that drive the LAN, minimize processing by running the storage devices off a stripped down specialized OS rather than a general purpose server, and spend bandwidth to preserve simplicity and an orientation to the Internet.

On this account, SANs should be on their way to extinction, their mazes of complexity an intolerable bottleneck in the face of network speeds that surpass computer I/O speeds.

So why are SANs still hanging around?

For the answer, consider an alternate history of the SAN, proposed not by SAN's defenders, but by the leading NAS insurgent Network Appliance itself, in the person of their most prolific and persuasive evangelist, Dave Hitz. In Hitz's view the SAN was invented in 1994 by the then leading independent Internet Service Provider. Swamped by the initial tsunami of Internet data, NetCom (NECS) kept multiplying servers with tethered RAID (redundant arrays of inexpensive disks) arrays to no avail. With the World Wide Web coming on like the Bay of Fundy, Netcom tried something new. It attached two "news" servers to a FDDI ring and on

the same ring attached two file servers to RAID drives. This was a breakthrough, because for the first time it separated Net servers from storage retrieval systems and allowed independent scalability of each. In theory, need more servers, add a server. Need more storage, add storage.

This SAN was contrived not to conserve network bandwidth (the whole point of the array was to get more data onto the Net faster) nor to augment storage space (counted in bytes per day, NetCom served orders of magnitude more data than it stored). The purpose was to address a processing crisis, the challenge of serving from a single source a vast array of data selections to a rapidly mounting mass of simultaneous users.

Although the field is crowded and there is scant room for mistakes, Storage Networks is the prime storewidth service provider today.

Pervading the Web today, that same processing crisis reflects not a shortage of bandwidth but its onrushing abundance. In the estimation of Ellen Hancock beyond light-speed latency, eighty percent of network delays are attributable to inefficiency at the nodes. Host computers on the Net are increasingly overwhelmed by the demands millions of networked clients can make on them.

On a white-board in the offices of the Gilder Technology Group, sits a diagram of what will be our third website architecture in less than a year. It was ably assembled by David Dortman when it became clear that Gildertech 2.0 would be recorded in history as a reconnaissance mission not unlike the scene in "Zulu" in which the tribal chief sends several hundred of his bravest warriors shouting and gesturing toward British riflemen in order to test the enemy's firepower.

At ten times the previous annual cost, the new site is elaborately configured for a single purpose—speeding I/O between the site and our subscribers. To achieve this goal, we do not move our data closer to the Net; we push it further away, behind serried ranks of laboring boxes, four ranks deep on the data center network itself (Cisco's (CSCO) pride and joy), through another rank of firewall boxes, thence to a set of load balancers. Only after leaving the load balancers, the sixth rank deep, do we hit the "Front Net," across which the load balancers distribute requests to two triads of web servers. Storage finally? Not yet. Should the Web servers need to plumb storage they must send requests through the "Back Net" to two clusters of SQL (structured query language) database servers, each sharing—through fibre channel links and across redundant hubs—a set of Storage Works RAID controllers, i.e. storage. The whole thing is preposterously complex, processor-ridden, power hungry, and apparently impossible to do without. Storage Networks—help!

Since EMC has been the titan of storage for the last

decade, everyone in the industry is currently targeting this \$8.6 billion behemoth for disruption. But most of the challengers to EMC fail to grasp a key part of Clayton Christensen's disruption paradigm. Disruption depends on the technology reaching "overshoot": the place on the learning curve where a technology has overrun the real needs of most customers.

But in the prevailing World Wide Web storewidth arena, undershoot is still the problem. The industry is nowhere near supplying fast and seamless access to diverse Web resources. Complex SANs are not throwing processing at an illusory bandwidth shortage. They are throwing networked processors at a very real processing crisis.

Taking a climactic step in the hollowing out of the computer, NetApp's Network Attached Storage pulled the filer operating sub-system—which makes and records decisions about where to place data on disk—out of the main server operating system. Rather than accessing the disk array as a physical set of devices organized in "blobs," the appliance addresses the actual files on the array no matter how they are physically organized. File-based disk management allows the NetApp file system, called WAFL (Write Anywhere File Layout), unprecedented freedom to optimize data location for later retrieval, substantially improving performance. More important, on the Net, it enables users to share files between Unix and NT, which is beyond the ken of any traditional bloc-based system, which treats a disk bloc as a black box. EMC is now making similar NAS appliances called Celerra.

Both companies, however, are responding to undershoot with optimized and highly integrated systems. NetApp's proprietary operating system, run on a speed freak Alpha processor, has little in common with the off-the-shelf engineering style of most "disruptive" products. Of what it estimates to be ultimately a \$30 billion storage subsystems market, NetApp believes that the bottom (and fastest growing) two-thirds can fall to NAS. To achieve this goal, however, even more specialized and dramatic performance tweaks are in the offing, further eliding the distinction between NAS and SAN. Whatever they are called, these devices will be prominent components in ever more ambitious attempts by high traffic, data dense nodes, such as SNI's, to share out data at Internet speed.

The market will be big enough for all. Of the some two hundred million disk drives sold last year only one million or so came through storage appliance and storage networking companies. Today on the Internet are some two hundred terabytes of data. Enterprise data may total around six hundred petabytes, more than three orders of magnitude more. It will mostly move to the Net.

The paradigmatic storewidth companies will be those that coast on the abundance of bandwidth rather than furiously coping. Richly paradigmatic solutions will emerge from the lambda network itself, as relatively low performance parallel processors connected to cheap, abundant, dynamic, and ever more fine-grained light-wave circuits hang "no waiting" signs outside today's most impacted web sites.

In the meantime, SNI has taken a step in the right direction precisely by choosing to be a service company. It wastes bandwidth in the form of dedicated lambdas to move the SAN-server-Fibre Channel-RAID-NAS tangle off your premises and onto its own. And it uses optical bandwidth to move SAN's most precious processors—the wet-ware of top dollar IT staff—off your payroll and onto SNI's. Defining a space and a strategy it joins the Telecom list this month.

Procom sweeps storage

Meanwhile, in the lower reaches of the storewidth market, where the ratio of data served to data stored is lower, and the processing bottleneck less acute, NAS still serves its original purpose: cheap, simple data sharing that takes advantage of the reversal in network and I/O speed advantages. Here, almost by default, companies can trade cheap bandwidth for expensive processing right now liberating storage from imperious general purpose processors exactly as NAS was meant to.

Created four years ago before the storewidth revolution, when bandwidth inside the computer was still at least ten times greater than bandwidth outside it, NetApp products are overdesigned for these scores of thousands of companies. By contrast, the modular Procom devices were designed over the last year. Procom uses the Linux operating system, which is already optimized for multiprocessors. NetApp uses a proprietary OS optimized for serial processors. Procom uses conventional Pentium processors. NetApp still employs leading edge Alpha chips to extort the utmost in execution speeds from a single processor and faces problems in adapting to the next Pentium generation. NetApp achieves fault tolerance through Compaq-Tandem's proprietary ServerNet. Procom obtains similar results from Ethernet links, while at the same time allowing for interoperability, an issue that has plagued EMC with its "all-EMC" installation of SANs. Procom's Ethernet connection enables scalability and provides built-in multi-leveled security. NetApp now faces the challenge of upgrading its systems for multiprocessing and new applications. Procom is preparing to ship its system to Hewlett Packard for use in HP Network Attached Storage devices. In another product to be introduced this summer, Procom is said to be adapting its systems with an open Posix interface from Bell Labs that allows ready customization of its appliances for such special purposes as email, news, Apache web service, and streaming media.

The drastic change in the networking environment gives Procom a huge opportunity to sweep into the 600 petabyte arena for storage below the commanding heights of Web Hosting centers and Forbes 500 enterprises. NetApp will win the largest margins, revenues, and dollar growth. But learning curves feed not on dollars but on units. Procom can expand unit sales more rapidly. With a faster accumulation of units, Procom's

costs will drop faster and its performance will become increasingly competitive. At some point over the next several years, the company (and its many imitators) have a disruptive opportunity to invade NetApp's space. With the bandwidth winds of the Telecom at its back, it richly deserves a place at the Telecom Table.

EMC has been the titan of storage for the last decade, everyone in the industry is currently targeting the \$8.6 billion behemoth for disruption.

Mango tango

Perhaps the most pleasant surprise on our trip came at the beginning in Chelmsford, Mass., before we so much as boarded an airplane. There we encountered the most immediately fetching product on our travels—Mango's Cachelink. Based on the caching and clustering and shared memory algorithms developed by Steve Frank for the late lamented Kendall Square Research's massively parallel supercomputer, Cachelink employs the free storage in browsers to accelerate web access in the enterprise. Any web page accessed by anyone in an enterprise cluster is made instantly available to the rest of the company. As the system expands it becomes more efficient because it embraces more web pages. Everyone needs this. We are currently installing it on all the computers in our office in hopes of accelerating our Web access by 50 to 100 percent. We will tell you how it goes.

George Gilder & Richard Vigilante
June 16, 2000

TELECOSM TECHNOLOGIES

ASCENDANT TECHNOLOGY	COMPANY (SYMBOL)	REFERENCE DATE / PRICE	MAY '00: MONTH END	52 WEEK RANGE	MARKET CAP	
WINGS OF LIGHT						
Wireless, Fiber Optic Telecom Chips, Equipment, Systems	Lucent (LU)	11/7/96	11 25/32	57 1/4	49 13/16 - 84 3/16	186.4B
Wave Division Multiplexing (WDM) Systems, Components	Ciena (CIEN)	10/9/98	8 9/16	119 11/16	26 13/16 - 189	16.9B
Wireless, Fiber Optic, Cable Equipment, Systems	Nortel (NT)	11/3/97	11 1/2†	53 7/8	17 15/16 - 72 1/16	148.4B
Optical Fiber, Photonic Components	Corning (GLW)	5/1/98	40 15/16	193	53 9/16 - 226 7/16	53.6B
Wave Division Multiplexing (WDM) Components	JDS Uniphase (JDSU)	6/27/97	3 5/8	88	16 3/16 - 153 3/8	68.8B
Adaptive Photonic Processors	Avanex (AVNX)	3/31/00	151 3/4	68	47 3/8 - 273 1/2	4.4B
All-Optical Cross-Connects	Agilent (A)	4/28/00	88 5/8	73 5/8	39 13/16 - 162	33.3B
THE LONGEST MILE						
Cable Modem Chipsets	Broadcom (BRCM)	4/17/98	6	130 1/16	43 1/8 - 253	28.0B
S-CDMA Cable Modems	Terayon (TERN)	12/3/98	15 13/16†	55 1/2	13 3/16 - 142 5/8	3.4B
Linear CDMA Power Amplifiers, Cable Modems	Conexant (CNXT)	3/31/99	13 27/32	37 5/8	18 15/16 - 132 1/2	8.2B
THE TETHERLESS TELECOSM						
Satellite Technology	Loral (LOR)	7/30/99	18 7/8	7 9/16	7 1/2 - 25 3/4	2.2B
Low Earth Orbit Satellite (LEOS) Wireless Transmission	Globalstar (GSTRF)	8/29/96	11 7/8	8 9/16	7 7/8 - 53 3/4	0.8B
Code Division Multiple Access (CDMA) Chips, Phones	Qualcomm (QCOM)	7/19/96	4 3/4	66 3/8	22 5/8 - 200	49.2B
Nationwide CDMA Wireless Network	Sprint (PCS)	12/3/98	7 3/16	55 1/8	21 3/4 - 66 15/16	50.5B
CDMA Handsets and Broadband Innovations	Motorola (MOT)	2/29/00	172	94 1/4	26 1/4 - 61 1/2	22.5B
THE GLOBAL NETWORK						
Broadband Fiber Network	Level 3 (LVL3)	4/3/98	31 1/4	76 5/16	45 1/4 - 132 1/4	27.9B
Broadband Fiber Network	Metromedia (MFNX)	9/30/99	12 1/4	30 15/16	10 9/16 - 51 7/8	16.9B
Submarine Fiber Optic Network	Global Crossing (GBLX)	10/30/98	14 13/16	25 1/16	20 1/4 - 61 13/16	20.8B
Broadband Fiber Network	Northeast Optic (NOPT)	6/30/99	15 1/16	34 3/16	14 - 159	0.6B
Telecommunications Networks, Internet Access	WorldCom (WCOM)	8/29/97	19 61/64	37 5/8	37 7/8 - 64 1/2	107.8B
CACHE AND CARRY						
Directory, Network Storage	Novell (NOVL)	11/30/99	19 1/2	8 5/16	7 7/8 - 44 9/16	2.7B
Java Programming Language, Internet Servers	Sun Microsystems (SUNW)	8/13/96	13 3/4	76 5/8	26 15/16 - 106 3/4	121.6B
Network Storage and Caching Solutions	Mirror Image (XLA)	1/31/00	29	25 9/16	3/4 - 112 1/2	2.7B
Disruptive Storewidth Appliances	Procom (PRCM)	5/31/00	25	25	5 3/8 - 89 3/4	284.9M
Remote Storewidth Services	Storage Networks (STOR)	<i>Initial Public Offering June 27, 2000</i>				
THE MICROCOSM						
Analog, Digital, and Mixed Signal Processors	Analog Devices (ADI)	7/31/97	11 3/16	77	19 1/8 - 94 11/16	27.2B
Silicon Germanium (SiGe) Based Photonic Devices	Applied Micro Circuits (AMCC)	7/31/98	5 43/64	99 1/4	14 1/16 - 158 7/8	11.8B
Programming Logic, SiGe, Single-Chip Systems	Atmel (ATML)	4/3/98	8 27/32	38 3/16	9 5/16 - 61 3/8	8.5B
Digital Video Codes	C-Cube (CUBED)	4/25/97	23	17 3/16	14 1/4 - 33 3/8	0.8B
Single-Chip ASIC Systems, CDMA Chip Sets	LSI Logic (LSI)	7/31/97	31 1/2	52 7/8	18 9/16 - 90 3/8	16.2B
Single-Chip Systems, Silicon Germanium (SiGe) Chips	National Semiconductor (NSM)	7/31/97	31 1/2	53 7/16	17 11/16 - 85 15/16	9.4B
Analog, Digital, and Mixed Signal Processors, Micromirrors	Texas Instruments (TXN)	11/7/96	5 15/16†	72 1/4	26 5/16 - 99 3/4	118.2B
Field Programmable Gate Arrays (FPGAs)	Xilinx (XLNX)	10/25/96	8 7/32	76 1/8	20 1/4 - 88 7/16	24.5B

ADDED TO THE LIST: PROCOM AND STORAGE NETWORKS

* INITIAL PUBLIC OFFERING † SPLIT ADJUSTED THIS ISSUE

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 and Forbes Inc.

Monument Mills • PO Box 660 • Housatonic, MA 01236
Tel: (413)274-3050 • Fax: (413)274-3031 • Email: info@gildertech.com

EDITOR: George Gilder

PUBLISHER: Richard Vigilante

TECHNOLOGY ANALYSTS: Charles Burger, Mary Collins, Bret Swanson

MANAGING EDITOR: Debi Kennedy

DESIGNER: Julie Ward

SUBSCRIPTION DIRECTOR: Rosaline Fernandes

GENERAL MANAGER, GILDERTECH.COM: David S. Dortman

PRESIDENT: Mark T. Ziebarth

**FOR SUBSCRIPTION
INFORMATION
TELEPHONE
TOLL FREE:**

(800) 292-4380

Website:

www.gildertech.com

Copyright © 2000, by
Gilder Publishing