

Big Money in Data Bottlenecks

1 July, 2011

Reading about the latest growth in digital data traffic and storage capacity is an easy way to make the head spin

Last year, for the first time, digital data creation surpassed one zettabyte - the equivalent of 1,000,000,000,000,000,000,000 individual bytes, according to research firm IDC. That's a more than five-fold increase from the levels of just four years earlier, estimated at 161 exabytes, which was itself enough bytes to contain about 3 million times the amount of information in all the books ever written.

The explosion in data, not surprisingly, is putting a heavy strain on the keepers of all those bytes. Nearly half of the large enterprises participating in a recent Gartner Group survey listed data growth among their top three IT concerns. The respondents ranked maintaining system performance and scalability as their second-biggest worry. And virtually all are expecting to invest in more efficient ways to store, retrieve, organize and secure their vast archives.

Among venture capitalists, it's a trend that hasn't passed unnoticed, and it appears to be a central thesis behind several very recent, relatively large funding rounds. In May and June, VCs invested more than \$140 million in at least seven companies developing data center-focused technologies, according to Thomson Reuters (publisher of VCJ). Several pitch themselves as players in the WAN optimization space, enabling more efficient transfers of data between distant servers. Others are developing memory technologies aimed at lowering the cost and improving the ability for enterprises to get quick access to their most crucial data.

Both are areas, VCs say, where the right solution can translate into unusually rapid revenue growth.

"What we're seeing is enterprises have a lot of capex being outlaid into data center building, and in each we're seeing a higher density of computer servers," says In Sik Rhee, a general partner at Rembrandt Venture Partners and a co-founder of data center automation pioneer Opsware, which was acquired by Hewlett-Packard in 2007 for \$1.6 billion.

"Plus, with the growth of data, the need to have a faster, more efficient way of moving data is becoming more of a pain point," he says.

The chance to alleviate that pain point, Rhee says, was a key reason the firm led a \$15 million Series B round in June for Infineta Systems, a San Jose, Calif.-based company that recently began marketing a hardware and software system for moving traffic between data centers. It's in a niche, Rhee says, where a lack of competitive offerings from incumbent technology providers means enterprises are open to patronizing a startup.

There Are Exits, Too

For those startups with the right product lineup, entrepreneurs are finding there's plenty of money available, from customers and investors.

That was certainly the case for Fusion-io, which develops flash-based memory platforms for data centers. The Salt Lake City-based company raised \$234 million in its June 9 IPO after hiking its offering price multiple times. Though shares were down nearly a third from their first-day trading highs, the company maintained a valuation of about \$1.7 billion in late June, resulting in still promising venture-scale returns for its largest backers. [New Enterprise Associates](#) holds a 33% post-IPO stake while [Lightspeed Venture Partners](#) has an 11% post-IPO stake. Previously, the company raised \$114 million in venture funding between 2008 and 2010, according to Thomson Reuters.

The Fusion-io offering also gave followers of the data center performance space a clear look at just how quickly sales can ramp. In the nine-months leading up to April, Fusion-io posted revenue of just over \$125 million, up from \$25 million a year ago. Over the same period, it narrowed its loss to \$1.2 million, from nearly \$20 million a year earlier. It may help that Fusion-io's largest customers were very fast-growing companies themselves. Facebook, at the time of the IPO, was Fusion-io's largest buyer and accounted for just under half of its revenue. Apple, through a reseller, generated one-fifth of Fusion-io's revenue in the first quarter.

Fusion-io's pitch to customers is that it helps clear an increasingly problematic data center chokepoint. Legacy datacenter storage architectures, it says, can't supply data quickly enough to fully utilize the processing capacity of today's servers.

This results in an increasing number of underutilized servers. It's sort of the data infrastructure equivalent of a short order cook unable to flip burgers fast enough to satiate the appetite of a ravenous diner.

While processing performance has doubled approximately every 18 months, the performance of storage infrastructure has not kept pace, and this increasing gap between processing and storage performance has created a growing problem that Fusion-io says it resolves by relocating critical, or active, data from centralized storage to the server where it is being processed.

Milpitas, Calif.-based Pliant Technologies, another startup focused on easing storage chokepoints, also generated a tidy return recently for venture backers. The 4-year-old company, which develops enterprise flash-based storage devices, sold to SanDisk in May for \$327 million plus equity incentives.

"Four years is a quick exit, but that's also a sign that the market's heated up," says Barry Eggers, managing director at Lightspeed, which was Pliant's largest shareholder. "There will be more activity like this."

Follow-on Rounds Abound

Lightspeed, even with two recent exits under its belt, still has a lot riding on its wager that demand for and innovation around faster data retrieval, particularly through flash memory technology, will keep gaining traction.

Although the firm made what Eggers says was a "very healthy venture return" on Pliant (and stands to make an even healthier one from Fusion-io), it still has bets on the table for several other companies tied to that trend. One is Nimble Storage, a developer of hybrid systems that combine flash with high capacity disk to enable cheaper backup and recovery. Another is IO Turbine, which focuses on input/output bottleneck issues by using flash technology directly in VMware server environments.

It's not a sector in which VCs are expecting exits on just a few million dollars in funding, particularly for startups developing hardware. Even the most capital efficient companies will generally require upwards of \$30 million over multiple rounds. Out of seven companies in the data center space that raised capital in May and June, for instance, the smallest round was \$14 million and the largest totaled \$40 million.

That largest round went to Violin Memory, a Mountain View, Calif.-based developer of scalable memory arrays and caching systems that shares several traits with Fusion-io. One is Don Basile, its CEO, who was formerly Fusion-io's chairman and CEO.

The other similarity is its rapidly rising valuation. The company's \$40 million funding in June came with a pre-money valuation of \$400 million. Violin - which Basile says is on track to surpass \$100 million in revenue this year - also competes in bringing flash memory to the data center.

Yet while a few sought after startups are getting premium valuations, the funding environment isn't exactly bubbly. Given how many VCs were burned by costly hardware infrastructure investments in the last boom, there's significant pressure on entrepreneurs to keep costs down.

"In terms of hardware, the market is there, but investors have been a lot more sensitive to initial capital requirements," Rhee says. After so many companies during the dot-com boom that built very sophisticated systems and failed, there's been "almost a knee jerk reaction to be more conservative," he says.

He adds that one of the draws of Infineta was that it developed its basic product on a comparatively efficient budget of \$15 million to \$20 million.

Still, for those able to keep capex levels in line with VC expectations, fundraising can be fairly straightforward. Raj Kanaya, CEO of Infineta, says the company closed its last round on schedule and within the planned range, focusing on a set of venture capitalists with deep backgrounds in infrastructure. The company plans to use the biggest chunk of money for rolling out sales and marketing.

The Future in Zettabytes

One encouraging factor for Kanaya is that public market investors are taking a keen interest in companies developing data center bottleneck solutions.

He points to RiverBed Technology, a developer of WAN optimization hardware and software whose shares have more than doubled over the past year, and F5 Networks, a developer of application delivery technology that's also posted huge stock price gains in the past couple of years. Both grew revenue by more than a third over the past year.

"From a Wall Street perspective, they look at that and say there aren't a lot of spaces that are growing at that double-digit rate with high profitability and a high barrier to entry," Kanaya says. "The financial markets are really rewarding the growth."

If forecasts for data output hold true, there should be plenty of demand ahead. After hitting about 1.2 zettabytes last year, global digital data output is supposed to go up to 35 zettabytes by 2020, according to IDC.

To put that number in perspective, linguist Mark Liberman calculated a few years ago that all human speech ever uttered, stored in digital format, would take up 42 zettabytes. Even with the prospect of much cheaper and more efficient storage and processing in the next decade, that's still a mind boggling amount of data to manage. Kanaya sees this as a bullish prospect.

"Fundamentally, the wind's at our back," he says. "The needs we're addressing are continuing to grow and grow."