

CASE STUDY:

As Membership Booms, Plaxo Finds New Storage Regimen a Key to Keeping Pace with Demand

With over 50 million members, pioneer of the smart address book adds new storage muscle to keep its development staff at peak efficiency

If membership is a solid indicator of corporate health, Plaxo is flourishing. The originator of the “smart address book” saw individual memberships grow 233% in just over four years – from 15 million to 50 million between late 2006 and today. With that growth has come greater speed and sophistication in how Plaxo develops new applications for its members, and in how it manages the ever-expanding needs of the development environment.

Plaxo’s growth brings with it new IT challenges: among them, long build times for new applications and the complexity of data storage in its virtualized networking environment. But that was before it selected the Nimble CS240 converged storage array. Today Plaxo uses two CS240 arrays, replicating data from each one to the other, with the benefits of protecting the company’s vital intellectual property and giving application developers a new tool for maximizing their productivity.

Managing a Virtualized Environment

Plaxo runs dozens of virtual machines in a VMware operating environment. For Ethan Erchinger, VP of Operations, the move to Nimble for data storage in a complex virtualized storage environment made practical sense, and within a matter of days proved his decision right.

“The architecture of the Nimble system just makes sense,” he said. Nimble uses solid-state drives (SSDs) to offload reads from disk, using the superior read IOPS and latency of SSDs to maximize performance, while maintaining long-term storage on high-density, ultra-low-cost SATA drives. “That means the unused data from each system stays on large disks, but active data is largely available on super-fast SSD. That’s what just makes sense to me.”

Storage Profile: Plaxo

Challenges

- Sluggish data transfers, resulting in downtime in Product Development
- Complexity of multiple virtual machines feeding a single storage array

Solution

- Deploy two Nimble CS240 converged storage arrays

Immediate Benefits

- Reduction of 50% in code build times
- 80% reduction in development downtime
- Failover times of 15 minutes
- Smoothing of I/O peaks in backend storage
- Significant cost reductions in product development

Long-Term Objective

- Scale up with Nimble as Plaxo breaks new membership barriers.

He adds that the Nimble system is well equipped for a development environment in which storage arrays can sit “on call,” only to be followed by sudden demand that creates large bursts of read and write operations.

“Long-standing complaints from development about long build times are legendary,” he says. “Before Nimble, if we had storage failures, our developers might just as well have sat on their hands while waiting for recovery of, say, 100 gigabits of data from the server. Now, simply recovering from a snapshot, or bringing those developers onto our replicated system, saves the company time and money.” Today the Nimble arrays provide Plaxo what Erchinger calls “replication on the fly.” “We use both of our Nimble arrays to serve data live, giving us great performance, and replicate that data to each other. That replication serves as ‘near-line backups.’”

The complaints from development are now a thing of the past, he says, allowing staff to concentrate on actual production issues rather than internal processes. “At the same time we’ve gained a bit of ‘future-proofness’ in our system. Future features will provide us expandability with minimal effort and no forklift upgrades.”

And how does the architecture of the Nimble system work for Plaxo? “SSDs fit well in our development environment because it’s so bursty. So, when developers need to run big compile jobs, they can do so in a very short period of time. Storage that leverages SSDs intelligently allows us to read and write that data lightning-fast and smooth out the peaks, yet not waste huge amounts of money on pure SSD storage.”

Reducing Code Build Times

According to Erchinger, improvements in the time required for each new code build have been remarkable. Plaxo’s build times have dropped by up to 50 percent – the actual number depending on how many builds are underway at the time.

“In the previous configuration, we used standalone servers, and recovery of a system meant copying up to 100 gigabytes of data over the network for use on another system.” The process could require anywhere from an hour if another system was readily available, to a day or two if it meant finding another system, with the additional risk that the new system would be subpar in performance or reliability. But today a system failure would mean a simple failover to the second Nimble system – allowing the entire process to complete in approximately 15 minutes.

“The net-net is that we are fully replicated now,” he said. “With the complaints behind us we can concentrate on production-level matters.”

Managing Backend Storage

Even as Plaxo moved to a virtualized environment, system performance remained a challenge. Some developers on older hardware were equipped with a paltry two single-core CPUs and four gigabytes of memory. Today developers all have access to four virtual CPUs and eight gigabytes of memory, an “overprovisioned” environment that allows Erchinger to breathe more easily. And

system performance is consistent across the board – a more democratic and performance-enhancing environment for developers.

Still, he says, backend storage has to be designed correctly or it will become a bottleneck to access by multiple virtual machines. “You’re really piling a lot more I/O into a single location,” he says, contrasting the new virtual architecture with that of legacy architectures involving physical machines equipped with their own hard drives.

According to Erchinger, Nimble eliminates the challenge of I/O peaks. “I/O peaks can occur if you have four or five virtual machines trying to push I/O onto a single backend storage device. We found that the Nimble system handles those peaks nicely – and does so while remaining efficient in power usage. Similar to VM environments, with Nimble we don’t need to deploy multiple overpowered storage systems to achieve our goals.”

Boosting Performance and Productivity

Because Plaxo developers compile hundreds of code runs through the day, and storing and retrieving that code is the company’s lifeblood, sheer storage performance, or IOPS, might very well be the critical factor in ensuring the performance of the company’s development group. Yes? Well, yes but with some additional qualification, says Erchinger.

“Performance is absolutely critical, and that’s all about getting the highest achievable IOPS for the dollar. The IOPS we’re achieving with the Nimble array has proven game-changing for our development team.”

Today, he says, developers are saving the company time and money because they turn out products much more rapidly. Because code compilation takes a fraction of the time it once did, idle time for developers has plummeted and productivity is at a new high.

“I don’t get questions any more about why a certain developer isn’t able to work,” says Erchinger. “When management is happy, my job becomes infinitely easier.”



About Plaxo

Comcast pioneered the "smart address book" and today hosts over 50 million address book accounts for Plaxo.com and for Comcast's SmartZone integrated communications center. In 2008, Plaxo was acquired by its largest customer, Comcast, and today operates as several distinct, entrepreneurial business units in our Silicon Valley office.

These three businesses include: Plaxo, focused on delivering on the vision of a truly smart, socially-aware, and pervasive address book; the Social Technologies Group (STG), created after the acquisition to lead address book, identity, and social strategies for all of Comcast; and Tunerfish, a new social entertainment internal startup, incubated in STG.