

Plug-and-play storage

May 13 2005

For something with no physical properties, information certainly takes up a lot of storage space. And that same intangible asset generates plenty of very real headaches. Businesses of every size know all too well about the difficulty of managing an amorphous resource that, worldwide, is growing by nearly 30 percent a year.

Not surprisingly, there's widespread demand for a safe, effective and relatively inexpensive way to house all that data – ideally, a solution that can expand in a snap as a company's information-storage needs multiply.

HP Labs researchers are working on an answer that could replace the inflexible storage systems of today with a modular system, designed to quickly accommodate a fast-growing organization's changing storage needs.

"Traditional storage systems are large, monolithic and expensive," says Alistair Veitch, a researcher in the HP Labs Storage Systems Department. "Every manufacturer has multiple product lines. They're targeted to different market segments and different specialties. Often, you've got separate teams designing separate product lines, and none of them use the same hardware and software."

As a result, when customers outgrow one storage system, they typically must start over as they move up to the next one.

HP already offers a relatively standardized approach provided via its next-generation storage architecture, the HP StorageWorks grid, which

uses some elements of the lab's modular, system

The ultimate goal: to give organizations the ability to focus on using the information they have stored, rather than on managing the storage itself.

Smart cells

Researchers began work on what became their modular system about three years ago.

"We looked at whether we could build a system that would scale from the low end to the high end and have all the features that you want," says Veitch, senior technical lead and project manager. "We asked whether there's some way we could use off-the-shelf hardware and software, rather than following the expensive, time-consuming tradition of customizing new technology for each new product line," Veitch continues. "And," he adds, "we looked at whether we could make that system for less money."

Based on the team's ongoing work on the project, the answer to all of those questions appears to be "yes." Essentially, their system relies on small, rack-mounted storage appliances consisting of disks, CPUs, storage adapters and network cards. Together, each of these constitutes one modular storage unit called a "brick" or "smart cell."

Cells can be added as needed. "By adding more cells, you get more capacity and performance," Veitch says.

Rapid response to demand

A software program "glues" cells together so they function like a single array, or, as Veitch puts it, "one very large pool of storage" that can be

allocated as needed. The storage load is automatically rebalanced whenever units are added or subtracted.

The result: a storage system that, unlike traditional disk-array storage systems, scales quickly and easily in response to demand. IT administrators can either automatically or manually deploy, expand or reconfigure their storage systems without disrupting service or affecting performance. And they can do so at relatively little expense because the use of common hardware components keeps costs down.

Upgrades of the product line are easy, notes Beth Keer, storage systems department manager. "You can adapt to changes in technology over time â€" improvements in interconnects, for instance â€" fairly rapidly."

Different, but effective

Researchers initially worried that their ideas might represent too dramatic a departure from traditional storage systems.

"We thought there might be a little psychological resistance to something that looks so foreign," Keer says. "But because of its scalability, it actually has a familiar feel."

Upgrades of the product line are easy, notes Beth Keer, storage systems department manager. "You can adapt to changes in technology over time â€" improvements in interconnects, for instance â€" fairly rapidly."

International, multi-disciplinary team

Like many teams at HP Labs, the group is both international and multi-disciplinary. Each of its nine members represents a different country and a variety of technical areas. (Veitch is from New Zealand; Keer was born

and raised in the United States and other team members hail from Argentina, Brazil, Denmark, Greece, India, Japan and Scotland.) Researchers' technical specialties encompass a lot of territory as well, and include distributed systems, availability, performance, algorithm design, modeling, engineering, code writing, operating systems, networking and storage.

That breadth of experience and knowledge helped them tackle the numerous technical challenges they faced in developing their solution.

One problem had to do with addressing component failures, which required replicating the data between cells. But incorporating backup protection into systems is a complex undertaking involving some tradeoffs, Veitch says. "Designing in hardware redundancy adds a lot to the costs and complicates the software."

Ultimately, the team struck the right balance, designing their solution to respond even better to failure than standard storage systems. When a piece of hardware on a traditional disk-array system fails, that system typically loses a large percentage — often half or more — of its performance capability. But with the HP Labs system, if one component fails, the others take over, allowing the system to function with little or no loss of performance. Administrators can also easily remove and replace malfunctioning hardware.

Different cells for different needs

The team built its prototype using standard rack-mounted servers. Because of the limited number of disks supported by these first systems, the team is now experimenting with newer designs that will incorporate larger numbers (up to 12 disks) in a single compact package.

One of the advantages of the experimental architecture is that various

aspects of the hardware are flexible â€” you can have a cell that has relatively more disk drives per CPU units, for instance, which would be more appropriate for bulk or archival storage. Other cells may have more memory for higher performance.

"What's exciting now is the huge opportunity that's ahead of us," she adds, "both in terms of the business engagements and all the research threads we can explore. It's probably more than we have the time and resources to do, so we need to prioritize. But that's a good position to be in."

If one fails, others carry on

Efficiently handling such heterogeneity is another key research area. Veitch anticipates that cells will come in a variety of storage sizes â€” from smaller, four-disk models containing a few hundred gigabytes at the entry level, to about a dozen or more drives containing up to a few terabytes.

To add additional storage, "You just plug in a cell and connect it to the network," Veitch says. "Then you go to your management console, where you're looking at the whole system. You'll see how just how much useable space is there."

The Labs experimental solution also uses different methods of replicating data to make sure it's always available, such as using an algorithm that provides mirroring of data in three locations. "That way, if one fails, the other two can carry on quite happily," Veitch says.

Continuing research

Although some HP StorageWorks grid products now use early versions

of smart-cell technology, the HP Labs team continues developing the concept.

Current experiments involve increasing system scale. "Right now, we can scale to 20 or 30 cells," Veitch says. "We'd like to go beyond that and build some really big systems."

How big? The team envisions architectures one day containing hundreds, perhaps thousands of cells – enough to build a virtual warehouse big enough to meet the storage needs of even the most prolific business.

Source: HP

Citation: Plug-and-play storage (2005, May 13) retrieved 22 September 2024 from <https://phys.org/news/2005-05-plug-and-play-storage.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.