

Managing the information that drives the enterprise

STORAGE

Vol. 6 No. 12 January 2008

SERIOUS STORAGE

Clemson University's James Bottum has big plans—5 petabytes of storage, high-performance computing and visions of IT prominence—for the South Carolina school. 20

ALSO INSIDE

State of storage standards **28**

Storage-hungry digital video **36**

Demystifying storage virtualization **44**


Clemson University bets on lots and lots of low-cost storage to enhance its profile and attract top faculty.
By Alan Radding



unlimited storage

EVERY COLLEGE AND UNIVERSITY YEARS to break into the top tier of the *U.S. News & World Report's* annual ranking of higher learning. Clemson University, a South Carolina state university that's probably best known for its nationally ranked football team, is no exception.






THE STORAGE INFRASTRUCTURE UPGRADE PLAN created by James Bottum, Mike Cannon and Jim Pepin (left to right) avoids big, costly enterprise storage arrays bundled with sophisticated firmware in favor of low-cost arrays that Clemson University will use as building blocks to assemble any kind of storage new faculty members need.



This year, *U.S. News & World Report* ranked Clemson in a four-way tie at 67th in the Top National Universities category for 2008, and tied it with another university at the 27th spot in the less-competitive Top 50 Public National Universities-Doctoral category.

“Clemson is dead serious about becoming a Top 20 school,” says James Bottum, Clemson’s CIO and vice provost for computing and IT. The usual way a college does this is to woo celebrity professors with big money or to build new student life centers with sleek facilities that dazzle students. Clemson has decided to substantially upgrade its data



“WHEN I CAME HERE, THE TWO BIGGEST WEAKNESSES WERE LACK OF STORAGE AND THE LACK OF A DATA WAREHOUSE,” SAYS JAMES BOTTUM, CLEMSON’S CIO AND VICE PROVOST FOR COMPUTING AND IT.

center by building out a totally new IT infrastructure along the lines of the National Science Foundation’s (NSF) Cyberinfrastructure initiative. By doing this, it expects to attract a new crop of young faculty by promising to deliver the storage, bandwidth and CPU resources needed to support their world-class research efforts.

“Cyberinfrastructure is the primary backbone that ties together innovation in research, instruction and service to elevate Clemson to the Top 20,” says Clemson’s provost and VP of academic affairs Doris Helms. She plans to connect Clemson’s Cyberinfrastructure to the NSF’s national research grid and lure 400 new faculty members to campus by 2010, when Bottum expects to have put in place 5 petabytes (PB) of flexible storage and, ultimately, 100 or more teraflops (TFLOPS, 1 trillion floating point operations per second) of processing. The school has 500TB of storage today and will have more than 1PB next year. “I’m confident we can get this done before all the new faculty are up and running,” says Bottum.

To get this done, the school projects it will spend \$3.5 million for storage replacement alone over the next five years. It also plans to request approximately \$10 million over the next five years to add to its base \$25 million IT budget. Finally, it’s upgrading the general campus network at a one-time expense of \$8 million (see “Clemson’s storage components,” on the next page).

Bottum, who was hired away from Purdue University specifically to lead this effort, and CTO Jim Pepin, another recent hire with this goal in mind, are leading an approach that doesn’t hesitate to break from conventional IT wisdom. Their storage infrastructure upgrade plan avoids big, costly enterprise storage arrays bundled with sophisticated firmware in favor of low-cost arrays (LCAs) that Clemson will use as building blocks to assemble any kind of storage new faculty members need.

And storage will be the lead component in Clemson’s Cyberinfrastructure push. “When I came here, the two biggest weaknesses were lack of storage and the lack of a data warehouse, a single authoritative data source,” says Bottum. First up: addressing the storage issue.


Storage condominium

To deliver the storage, Clemson will develop what Pepin, who was lured from the University of Southern California (USC), calls a “storage condominium cluster.” The storage condo is assembled out of blocks of LCAs. When a faculty member requests storage for his or her group, it will be assigned a set of LCAs as its storage condo. The group will then be responsible for provisioning, configuring and managing that particular storage however it pleases. The Clemson storage team will maintain the infrastructure and be responsible for overall infrastructure performance and reliability, but not the individual storage condos.

Clemson's storage components

Mike Cannon, data storage architect and manager, was brought in by Botum—for whom Cannon had worked at Purdue University—to lead the storage effort and build out Pepin's storage vision. The LCA, Cannon's core building block, currently consists of a tray of 12 300GB SATA disks with two controllers. Clemson is buying the LCAs initially from Sun Microsystems Inc. (list price is \$23,000, although Clemson negotiated a steep discount). Although the initial set of arrays is from Sun, it's just commodity storage. "We could get the LCAs from any vendor, and we are talking to others," says Cannon.

The LCA-based storage condo, however, is a far cry from what Cannon thought he was being hired for. "When we hired Mike [before Pepin arrived] we were thinking virtualization," says Boyd Wilson, executive director, computing systems and operations, and Cannon's immediate supervisor. Previously, Clemson's IT infrastructure resembled a hodgepodge of disparate systems that operated as



▶ **LCA UNITS:** Multiple trays of 12 300GB SATA disks with dual controllers

▶ **SWITCH:** QLogic SANbox 9200 switch, 16-port 4Gb/sec blades, dual CPU

▶ **REDUNDANT FIBRE CHANNEL NETWORK**

▶ **TAPE BACKUP:** 3,000-tape Sun Microsystems StorageTek tape library

▶ **MIRRORED DATA CENTERS:** 15km apart

highly independent silos.

Pepin arrived soon after Cannon and convinced everyone that virtualization was too complicated and costly, not only in terms of money but in performance and bottlenecks. "We were doing this massive upgrade to eliminate bottlenecks. When Jim [Pepin] came, he said we don't need virtualization, we need to get simple," says Wilson. "That's when he started talking about this idea of LCAs. I was petrified. I thought it was too small."

It was new to Cannon, too, but he was willing to try. It turns out that "when we need storage or more IOPS, we can sum the LCAs to meet the need," he says. That was the beauty of Pepin's storage condo. If a group needed 50TB for a massive research project, they could assemble it. If another group needed 2TB for some IO-intensive database app, they could give them that, too. Each need would be configured, provisioned and managed by the users to their own specs.

The trick turned out to be using low-cost switches in conjunction with LCAs. The switches perform the same job in the network layer that a virtualization engine would otherwise do. "We're not afraid to carve up the storage and rely on the switches to connect it together right," says Cannon. Clemson currently uses four QLogic Corp. SANbox 9200 switches with eight 16-port 4Gb/sec blades, dual CPU, but no dual licensing for fault tolerance. The street price for the switch is approximately \$700 per port. Through negotiation and by figuring in additional discounts in the form of grants, Clemson paid somewhat less.

Before he can carve the storage into condos, however, Cannon has to do considerable analysis. "We

TO DELIVER THE STORAGE, CLEMSON WILL DEVELOP WHAT CTO JIM PEPIN CALLS "A STORAGE CONDOMINIUM CLUSTER."

have to look at the application and see how it reads and writes data before we allocate the storage,” he explains. In the end, the team is getting to the same point as it would with virtualization but getting there faster and at a lower cost. But there’s one tradeoff: “With virtualization maybe we would get better utilization,” says Cannon.

Cyberinfrastructure

The storage condo is just one piece of the Cyberinfrastructure that Clemson is rushing to build. The overall Cyberinfrastructure will encompass data and voice communications, systems and operations. For

users, it will include enterprise applications, learning and collaborative technology like Google Apps, and research and scholarship capabilities. It will support high-performance computing and provide additional access to pools of idle processing power through a project called Condor, which scavenges idle cycles for various learning and research projects.

“Condor is a scheduling and computational system that allows researchers to use extra cycles on desktop lab computers to work on scientific and other computational problems,” says Clemson’s Wilson.

The results will eventually be stored in the researcher’s storage condo. Clemson currently has more than 1,000 workstations in the Condor pool.

Supporting all of this will be a dual-redundant IT infrastructure (still under construction) consisting of two data centers, each built for full redundancy. Each data center, 15 km apart, contains a dual Fibre Channel (FC) SAN. Using a volume mirror, data is replicated between sites. Data will also be backed up to tape using a 3,000-tape Sun high-end tape library.

Hosts run the Sun SAM-QFS clustered file sys-

tem, which lets multiple hosts access the same volume, reading and writing simultaneously. SAM-QFS also sends the data to tapes automatically. Although Clemson’s legacy Legato backup system is no longer needed to write server data to tape, it’s being used to move data to SAM-QFS.

LCAs and some pieces of the new infrastructure have already been installed and are in production. Other pieces, such as the tape library, have been ordered but haven’t arrived. The SAM-QFS component is not yet working “but Jim Pepin has already done it at USC, so we’re confident it will work,” says Cannon.

The LCAs in the completely redundant FC fabric are attached to the QLogic switches. With two separate dual-FC fabrics, “there is no single point of failure. We have redundancy for everything,” says Cannon.

But there are limits to everything and Clemson didn’t use redundant switches within each fabric. “We are not building redundant redundancy,” says Cannon. “The cost of the dual switch licensing is expensive.” Dual switch licenses would cost Clemson \$50,000 for each license, and \$200,000 for the dual pairs. Instead, the operations team will rely on an alternate path in case of a failure and initiate a manual failover process in the event of a problem. They’re betting that the likelihood of such a failure at an inopportune moment would be so rare that they could safely avoid the added expense.

Open systems and mainframe

The LCAs currently work with Clemson’s open-systems hosts. The school’s numerous hosts run a variety of OSes, including Solaris, SUSE Linux, Red Hat Linux, NetWare and Windows. The systems support student, administrative and faculty applications. Microsoft Exchange is a big application. Individual user storage quotas are being increased from 50MB to 2GB.

In addition, the storage condo will connect to the back end of a separate high-performance computing cluster that currently provides 11 TFLOPs and expects to reach 30 TFLOPs soon. Bottum plans to push that up to 100 TFLOPs or more. “The plan is to use commodity-based storage and networking to support all of this,” says Cannon. Clemson is wooing new faculty with the promise of 140TB of high-

NSF Cyberinfrastructure

The goal of the Cyberinfrastructure is to facilitate new apps, collaboration and interoperability across institutions and disciplines. It encompasses the following:

- Computing cycles
- Broadband networking
- Massive storage
- Managed information
- Shared standards
- Middleware
- Basic apps for scientific computation

Source: National Science Foundation (NSF) Cyberinfrastructure Report

performance disk immediately at 10GigE interconnect speeds.

The school has an old EMC Corp. storage system that is at end-of-life and won't be continued. Also going away are the storage silos that characterized Clemson's IT environment in the past. A single open-systems storage team under Cannon is now pulling in all the various silos. "It's exciting. It's more stressed, but more fun," he says, speaking for the team.

A small mainframe group manages an IBM Corp. System z800 running z/OS. The group runs



“WE’RE NOT AFRAID TO CARVE UP THE STORAGE AND RELY ON THE SWITCHES TO CONNECT IT TOGETHER RIGHT,” SAYS MIKE CANNON, DATA STORAGE ARCHITECT AND MANAGER.

the university's core student record application but its primary workload is running the state's Medicaid app. The school has allocated 2.4TB of storage to the System z.

“With the student data and the Medicaid application, the mainframe is not going away,” says Cannon, although the primary thrust of the Cyberinfrastructure initiative is away from the legacy environment (see “NSF Cyberinfrastructure,” previous page).

New data center

The showcase of the new Cyberinfrastructure is a 30,000 sq. ft., \$25 million primary data center. It sports new power supplies, new UPS systems, dual/redundant everything, extensive FC wiring, and racks of servers and storage laid out along hot and cold aisles. Extra power, cooling and wiring capacity to support the most optimistic growth the university's provost ever envisioned are all incorporated in the design.

At the same time, Clemson has solved its power issues for some time to come. “We negotiated a deal with Duke Power [now called Duke Energy Corp.] to run our power and cooling and get us out of the power business. Duke can take us to 8Mw [megawatts],” says Bottum. Twenty TFLOPS, for comparison, draws just 130 kilowatts.

The new data center, situated on a wooded site that's part of a larger advanced research park Clemson is building a few miles from the main campus, was still under construction and only partially op-

erational when *Storage* visited. The centerpiece of the data center is the new network operations center (NOC). A far cry from the typical crowded NOC buried in a cramped basement and crammed with mismatched monitors and rejects at their final stop before the dump, the new glass-walled Clemson NOC is roomy, sparkling clean and features a row of workstations in front of a full wall of vivid, large-screen flat displays. Colorful graphics show the status of systems and networks, while a continuous CNN video feed demonstrates that the Internet connection is up and running.

Change management

As welcome as many of these changes are, they still upset well-established routines. Wilson and Cannon addressed training and change management from the outset. “We had a lot of independent silos,” says Cannon. “People didn't talk much. In many cases we were asking people to learn new skills.” Pepin's idea of the storage condo made of LCAs connected by switches was new to everybody.

Change management revolved primarily around meetings in which Cannon would explain what was happening and how the team would proceed. “We were going to more meetings, but now things were getting done,” says one team member. But resistance was expected. “We have some EMC users who need to be sold away from EMC,” says Cannon. “They need to learn new skills and need new documentation.”

To keep everyone up to date, Cannon created a wiki that contains the provisioning and configuration documentation of all the storage. Every team member goes to the wiki first whenever they have to do something with the storage and whenever they change something. Click on any component and the team member can drill down to more detail. "This lets us logically build the array. We don't even have to be here physically to do the work. We can do this over the network," says Cannon. A separate hardware group handles whatever actual physical work is required.

Barriers

Bottum exudes confidence in the Clemson Cyberinfrastructure initiative. "I'm sure we can complete it. We have a business plan with a technology strategy. We know what it's going to cost," he says. Speed, however, may prove the biggest challenge. "I worry about the pace. We're making big changes while keeping things going. HPC [high-performance computing] is a challenge, but so is email," says Bottum.

Another worry, but not Bottum's, is the ability to hire 400 top-notch faculty members in a few years.

What Bottum does worry about is managing the data the new faculty will generate. "The next challenge is the lifecycle management of data repositories," he says.

The biggest barrier may be unspoken but obvious—the political will at the state level. Clemson is a state university and is therefore subject to the politics that affect any state organization, especially budget politics.

Will storage condos, 100 TFLOPS HPC server farms, thousands of workstations in Condor processing pools, high reliability, seemingly unlimited storage and high-speed connections to the NSF research network be enough to catapult Clemson higher in the *U.S. News & World Report* college ranking? Provost Helms is counting on it. Bottum and Pepin haven't said so explicitly, but many on the staff expect them to finish their careers on an up note at Clemson. It sounds like a solid plan, but a nationally ranked football team might also play a big role. ☉

Alan Radding is a frequent contributor to Storage.