

How Boston University is making the most of its IBM zSeries and TotalStorage environment

BY JIM UTSLER



# A LESSON

**W**ay back when I was in college (in a galaxy far, far away), computers and most things related to them were something of an oddity, especially for a liberal arts major such as myself. We didn't use e-mail, the Internet was still a green-screen mystery and I did all of my class research in a library, using—if can you believe it—an alphabetized card catalog.

Things have changed quite a bit since then (even though this was less than two decades ago), and such things as e-mail and online research are the norm, making pertinent communications and information gathering a snap. Of course, with these blessings also come some overhead, not the least of which are increased needs for both computing power and data storage, especially as colleges and universities continue to encourage faculty, staff and students to take advantage of the online environment.

The challenge becomes finding the right combination of reliable hardware and software to lessen the burden on the IT staffs of

institutions of higher education. If not, they risk possible meltdowns and mounds of overtime, with IT personnel working late into the night to put out fires and get everything back online in time for peak computing periods.

One organization that has largely met these standards is the Boston University Administrative Computing Services (ACS) datacenter. Using an IBM<sup>®</sup> zSeries<sup>®</sup> server, IBM TotalStorage<sup>®</sup> solutions and applications such as FDR/UPSTREAM from Innovation Data Processing, it has created a rock-solid IT environment that allows it to easily meet its users' needs, letting the

# LEARNED

aforementioned faculty, staff and students conduct much of their university-related business online and with confidence.

## Core Classes

Located on the banks of the Charles River in Boston's historic Back Bay district, Boston University has been in existence for some 157 years, beginning as a training institute for Methodist ministers. Since then, it has become the fourth-largest independent university in the United States, having more than 29,000 students representing 143 countries and all 50 U.S. states.

Because of its large size, Boston University has rigorous IT requirements, including basic back-office functions such as HR, payroll and other financial tasks. And because of its

capabilities are only one reason the university, which has been a mainframe shop for nearly 30 years, brought the new box into its IT environment in June 2003. "We appreciate its workload management and data sharing, and its high reliability, availability and serviceability, to name just a few things off the top of my head." The mainframe offers the university an economy of scale not available with other platforms. The university is also considering the use of the server's Capacity on Demand (CoD) capabilities "when there are un-forecasted spikes in business need, although we do a fairly good job of projecting computing requirements," Shockley adds.

Shockley estimates that most of the university's data load is read rather than write, pegging the numbers at around 90-

**"It's a highly reliable disk storage unit that's simply just there. It's never down, and if it hiccups, a service technician shows up with a part and we schedule a dynamic installation."**

**—Gerard Shockley, assistant director of technical services, Boston University**

position as an educational institution, it has additional unique needs, such as student and faculty support, including, according to Gerard Shockley, assistant director of technical services with Boston University, "classroom scheduling, student/faculty advising and degree auditing, to name a few."

In short, the organization relies heavily on its IT underpinnings, wanting to make sure its 60,000 or so customers, who include faculty, staff, students and alumni, are well served. To that end, it has a zSeries 890, several AIX\* technology-based servers and more than 20 Windows\* servers. (Despite this somewhat heterogeneous server environment, Shockley stresses that the university puts its "chips on the zSeries server.") It also has a TotalStorage Enterprise Storage Server\* (ESS) F20 with 2.1 TB of storage and two IBM 3494 tape libraries with 3590-E tape drives, with each of its tapes capable of holding as much as 120 GB of data.

Several virtual Linux\* partitions operate on the z890, including one that runs a Java\* administrative browser that allows users to view their mainframe reports through the Internet as soon as they're ready. As Shockley explains, "Instead of using a 3270 application, users can tap into the Java interface via a browser and customize columns and views, export data in a graphical format and translate the data to a number of different formats so they can massage it. This is a very popular tool." Notably, the university is strongly considering a road map to migrate its AIX applications to Linux on zSeries where it makes good business sense.

Shockley notes, however, that the zSeries virtualization

percent read and 10-percent write. As he notes, "We allow global access and the z890 is Web enabled." To support these massive data needs, the organization has put the ESS F20 in place to act as its data "backbone," Shockley says. "It's a highly reliable disk-storage unit that's simply just there. It's never down, and if it hiccups, a service technician shows up with a part and we schedule a dynamic installation."

## Declaring a Major

Because its data is vital to everyday operations, Boston University has put great effort into making sure the data is available around the clock. To that end, it backs up its data, whether from the mainframe, the mainframe's virtual Linux partitions or the distributed servers, to the F20 for quicker data access and to the tape for archival and disaster-recovery purposes. The key, however, was to find a way to move all of its disparate data, whether from server DASD or the F20, to tape for backup and recovery.

The mainframe backups are handled by Innovation's FDR product, which, as Shockley explains, "does the mainframe key data-volume backup." However, different methods were needed to back up the distributed systems, including the AIX and Windows servers and the virtual Linux on zSeries servers on the z890. In those cases, the organization uses Innovation's FDR/UPSTREAM product, which "does Linux on zSeries backups and restores, as well as the distributed file backup and restores for the distributed data," Shockley says.

The mainframe acts as the central administration point for

**CUSTOMER:** Boston University**LOCATION:** Boston**BUSINESS:** Higher education**HARDWARE:** One IBM <sup>^</sup> zSeries 890 and one IBM TotalStorage Enterprise Storage Server F20**CHALLENGE:** Backing up both its mainframe and distributed servers**SOLUTION:** Using Innovation Data Processing's FDR and FDR/UPSTREAM to create a largely automated data-backup process

these distributed backups, which include full-volume mainframe backups using FDR and incremental backups using FDR/UPSTREAM. (Full Linux on zSeries backups are also accomplished using FDR/UPSTREAM.) According to Dave Palermo, a project leader with the university's storage-management group, "FDR/UPSTREAM is a dual-component system. There's a piece that resides on the mainframe under VTAM as a started task with multiple sub-tasks, and then there's a piece that resides on the clients, on the Windows and AIX file servers. The way it's set up here is that MVS batch jobs are initiated and use TCP/IP to connect to the client and, based on file specifications that are in the batch jobs, execute, connecting with the client system and then performing the specified type of backup. It then goes across the network to transmit the backup data from the servers to our mainframe storage."

The university is also using an FDR/Upstream feature called full-merge backup. This allows the organization to perform a full backup and then run a series of incremental backups, which can be stored to tape. Those changes are then consolidated or merged with the full backup of that server. Christine Ciocca, storage-management technician with Boston University, explains, "This feature will look at the server, send a directory listing up from the client, compare it against the mainframe database to see what's been backed up, and then, if a file has been changed or added, request that that file be sent up. Then, on the mainframe, those incremental changes will be merged with the full, giving you an exact image of the server at that point in time without having to deal with all of the unchanged files. This saves a great deal of time because you don't have to perform multiple full backups."


The decision of whether to perform full or incremental backups is made by the individual server administrators, depending, for example, on the criticality of the server function, server recovery requirements and the amount of data they expect to change. The more critical the data is, the more likely it is to undergo a daily incremental with a weekly full-merge backup. If the data isn't as critical, a weekly full backup will serve the purpose.

The university has used other products in the past to attempt

to achieve the same results, but it found them wanting, both in features and in terms of the amount of manual labor involved. For example, prior distributed-system backup systems required someone to mount and dismount tapes, not only a time- and resource-wasting effort, but also impractical in an unattended computing environment. Currently, with its automated IBM tape libraries and FDR/UPSTREAM, this and other manual chores have essentially been eliminated. "Now," Shockley says, "I have two people, and they never have to leave their chairs to initiate a backup or restore."

This shift in backup paradigm has allowed storage administrators to become "engineers," as Shockley puts it. Automation of the backup processes means they can now contribute more fully to the organization, being repurposed to other, more interesting, more meaningful tasks. "We've actually created a different shop here, with IT staff members now becoming much more productive," Shockley points out. "For example, they can now construct rules of new data using storage-management utilities, design, implement and verify disaster recovery and business resumption, and review performance and capacity planning of storage across all platforms. We're now much more proactive than reactive."

### Learning from the Best

Back when I was in college, many of the computing capabilities mentioned here hadn't even been considered. Or, if they had, the execution wasn't nearly as seamless as it is today. As the case of Boston University shows, more data doesn't necessarily mean more work. In fact, its automated IT environment, with its IBM zSeries server, IBM TotalStorage technology and software from the likes of Innovation, stands as a lesson to other organizations about how best to handle large computing and data volumes, whether in the private or public sector. 

**Jim Utsler,** <sup>^</sup> Magazine senior writer, has been covering the technology beat for several years. His responsibilities include writing case studies and other articles of general computing interest. Jim can be reached at [jjutsler@provide.net](mailto:jjutsler@provide.net).