

Data Center Consolidation at the University at Albany

Rebecca L. Mugridge
and Michael Sweeney

ABSTRACT

This paper describes the experience of the University at Albany (UAlbany) Libraries' migration to a centralized University data center. Following an introduction to the environment at UAlbany, the authors discuss the advantages of data center consolidation. Lessons learned from the project include the need to participate in the planning process, review migration schedules carefully, clarify costs of centralization, agree on a service level agreement, communicate plans to customers, and leverage economies of scale.

INTRODUCTION

Data centers are facilities that house servers and related equipment and systems. They are distinct from data repositories, which collect various forms of research data, although some data repositories are occasionally called data centers. Many colleges and universities have data centers or server rooms distributed across one or more campuses, as does the University at Albany (UAlbany). This paper reports on the experiences of the Libraries at UAlbany as the Libraries' application and storage servers were consolidated into a new, state-of-the-art, university data center in a new building on campus. The authors discuss the advantages of consolidation, the planning process for the actual move, and lessons learned from the migration.

BACKGROUND

The University at Albany is one of four university centers that are part of the State University of New York (SUNY) system. Founded in 1844, UAlbany has approximately 13,000 undergraduates, 4,500 graduate students, and more than 1,000 faculty members. It offers 118 undergraduate majors and minors, and 138 master's, doctoral, and certificate programs. UAlbany resides on three campuses: Uptown (the main campus), Downtown, and East.¹ The Uptown campus was built in the 1960s on grounds formerly owned by the Albany Country Club. The campus was designed by noted architect Edward Durell Stone in 1962–63 and was built in 1963–64. The campus buildings include four residential quadrangles surrounding a central "Academic Podium" consisting of thirteen three-story buildings connected on the surface by an overhanging canopy and below ground by a maze of tunnels and offices. Many of the university's classrooms, lecture halls, academic and operational offices, and infrastructure are housed within the podium on the basement or subbasement levels. This includes the university's original data center, which is located in a basement room in the center of the podium.

Rebecca L. Mugridge (rmugridge@albany.edu) is Interim Dean and Director and Associate Director for Technical Services and Library Systems, and **Michael Sweeney** (msweeney2@albany.edu) is Head, Library Systems Department, University Libraries, University at Albany, Albany, New York.



While visually striking and unique, the architectural design of the podium presented many challenges since its construction, one of which is regular flooding of the basement and sub-basement levels. The original data center was flooded many times, to the extent that any heavy rainstorm had the potential to disrupt functionality and connectivity. When the university was first built in the 1960s it was not known to what extent computing would become part of the university's infrastructure, and the room that the data center was housed in was not built to today's standards for environmental control, such as the need for cooling. At the same time, server rooms sprouted all over the university, with many of the colleges and other units purchasing servers and maintaining server rooms in less than ideal conditions. These included server rooms in the College of Arts and Sciences, the School of Business, the Athletics Department, the University Libraries, and many other units.

University Libraries' Server Room

The University Libraries maintained its own server room with two racks full of equipment that supported all of the Libraries' computing needs. These servers supported our website, MSSQL and MySQL databases, EZProxy, ILLIAD (interlibrary loan service), Ares (electronic reserve service), and our search engine appliance (Google Mini). They also included our domain controller, intranet, and several servers used for backup. Two servers and a storage area network housed our virtual environment, containing an additional nine virtual servers. These included servers to support library blogs, wikis, file storage, development and test servers, and additional backup servers. The only library servers not housed in the Libraries' server room were the integrated library system (ILS) servers that were maintained primarily by the university's Information Technology Services (ITS) staff, our backup domain controllers, and a server holding backups of our virtual servers. The ILS production server was housed in UAlbany's data center and the ILS test/backup server was housed in the alternate data center in another building on campus. Also, two of the Libraries' backup servers for other applications were housed in the university data center.

The Libraries' server room consisted of a 340 square foot room on the third floor of the main campus library that was networked to support servers housed in two racks protected by a fire suppression system. There were two ceiling-mounted air conditioning units that cooled the room sufficiently for optimum performance. The Libraries' Windows System Administrator's office was nearby and had a connecting door to the server room, giving him ready access to the servers when needed.

Data Center Consolidation

Data center consolidation is defined as "an organization's strategy to reduce IT assets by using more efficient technologies. Some of the consolidation technologies used in data centers today include server virtualization, storage virtualization, replacing mainframes with smaller blade server systems, cloud computing, better capacity planning and using tools for process automation."² In addition to the investigation and use of these technologies, the planning for a

new data center often involves the construction of a new building or the renovation of a current building.

There were several drivers behind the UAlbany's decisions to build a new data center. In addition to the concerns mentioned above about the potential flooding risk of the current data center, the ability to manage optimum temperature was also a factor. The current data center was built to house 1960s-era equipment and was not able to keep up with the cooling requirements of the more extensive computing equipment in use in the twenty-first century. The current data center also occupied what is considered prime real estate at the university, at the center of campus and near the Lecture Center, which experiences high foot traffic during the academic year. The new data center was constructed near the edge of campus, with little foot or auto traffic, allowing the space previously occupied by equipment to be repurposed in a way that better meets the university's needs.

Like many other universities, UAlbany is increasingly making use of cloud computing capabilities. For example, the email and calendaring system are cloud-based. Nevertheless, this movement is being made in a deliberate and thoughtful way, leaving many of our administrative computing needs reliant on the use of physical servers. UAlbany and the Libraries have decreased the number of physical servers necessary by relying on a virtualized environment, and part of the project to move to the new data center included a conversion from physical to virtual servers. The Libraries' ILS production and test servers remain physical, as do several of the other Libraries' application servers. Many of the Libraries' backup servers are now virtual.

While there was no official mandate to consolidate all of the distributed server rooms across campus into the new data center, everyone involved understood that this was a direction the university administration supported. The Libraries' dean and director also supported this effort on behalf of the Libraries and charged Libraries' staff to collaborate with ITS to make this happen. Some of the drivers behind this decision include the promise of a better environment, improved security, backup generators for computing equipment, the use of ITS's virtual environment, the automation of server management, a faster network, the ability to repurpose the Libraries' server room, and more. These drivers are described in more detail later in this paper.

Construction planning for UAlbany's new data center began in the mid-2000s and included the identification of funding and the architectural design of the new building, later to be named the Information Technology Building (ITB). The actual construction began in 2013, with an estimated completion date of February 2014 and occupancy in April 2014. Unexpected challenges during construction delayed the timeline somewhat, and the construction was not completed until May 2014. The certificate of occupancy was granted in fall 2014. The data center is certified as Tier III by Uptime Institute,³ and the building is designated LEED Gold.

Alternate Data Center

Simultaneously with the construction of the new data center, the university entered into an agreement with another SUNY institution to house our alternate data center. This center was originally housed in another building on the UAlbany campus, less than a mile from both the main data center and ITB, in a building leased by UAlbany. This situation left some environmental issues out of our control, not an ideal situation. For example, an air conditioner failure in fall 2013 caused our backup and test ILS server to be down for six days, affecting our ability to use that server for other purposes and holding up several projects. In addition, data center best practice calls for an alternate data center to be housed at a distance from the main data center. In February of 2014, the servers in the alternate data center were moved to their new location. This included the Libraries' backup ILS server as well as two backup servers formerly housed in the main data center.

Advantages to the Libraries Moving to the University Data Center

There were many advantages to the Libraries moving to a centralized data center. Many of these advantages also applied to the other units considering a move to the new data center, but for the purposes of this paper, we are addressing them in the context of the Libraries' experience.

Repurpose Space

The Libraries' server room occupied a large office that could be repurposed to house multiple staff offices or student spaces. The Libraries have many group study rooms available for student use; however, they are in great demand, and the possibility of gaining more space for student use was seen as an advantage to making the move to a new data center.

Climate Control

The new data center is built on a raised floor that allows better air circulation. Hundreds of servers and other pieces of equipment create a lot of excess heat, and raised floor construction allows for better circulation of air. New racks have chimneys that exhaust heat from high-density computing environments. Air conditioners supply a constant stream of air that will maintain the optimum temperature for computing equipment. Sensors continually monitor humidity and keep it at an optimal level. This was an improvement over the Libraries' current server room, which had sufficient air conditioning for our relatively small number of physical servers but did not have backup generators to keep equipment running during a power outage.

Backup Generators

The new data center was built with two backup generators. If the building suddenly loses power, the backup generators will immediately start and provide a seamless source of energy. A secondary benefit to the university is that the backup generators can also provide a source of energy to other buildings on that side of campus; this area did not previously have a backup source of energy. Again, the Libraries' server room did not have a redundant electrical supply. In

the event of a power outage, battery units would allow the servers to shut down properly if the outage lasted more than forty-five minutes.

Security

With server rooms scattered all over the university, security issues were a concern. Now that the servers are housed in one location, the university can provide a highly secure environment in a more cost effective way. The new data center has card-swipe access to the building and biometric access to the data center itself. There are also cameras installed in the building as a further security measure.

Virtual Environment

Although the Libraries have made strides toward moving into a virtualized environment in the past few years, we had many constraints on our ability to keep up with developments. The Libraries' virtual environment was two versions behind UAlbany's virtual environment, and the storage needs of the Libraries' virtual environment were at capacity. Part of the incentive to moving into the new data center was the ability to downsize some of our physical equipment and migrate some of our physical servers to virtual equivalents.

Automation of Server Management

One of the benefits of consolidating servers into one environment is that they are in a secure location, but it is still possible to manage them from a distance. The virtual environment has a web-based console that allows system administrators to connect and manage them, and the physical servers can be managed over the network as well. Even though the servers are centralized, our system administrator can work from an office in the library, or from home if needed.

Faster Network

Part of the project to construct a new data center included the installation of an additional fiber network across campus. The new fiber network connects all buildings on campus with each other and the new data center. All of the network equipment was upgraded, providing faster connections and response time. The additional fiber network is fault tolerant: if the primary network fails, the second fiber network can immediately take its place with no loss of service.

Staging and Work Room

The new data center was designed to include a staging and work room. This can be used by any of the system administrators who are responsible for equipment housed in the data center, and it allows them to work on equipment in a room adjacent to the locked and secure data center.

Equipment Inventory

Part of the planning for the migration involved creating a detailed inventory of equipment. The Libraries already had a server inventory, but the information collected for the migration went far beyond just a list of the servers. This helped us identify who was responsible for physical and virtual servers and who was responsible for the services and applications that ran on those servers. Creating the equipment inventory also allowed us to consolidate and decommission equipment that was no longer needed, and additionally helped us determine a prioritization and time line for the move.

Applications Inventory

In addition to creating an equipment inventory, the Libraries created an applications inventory that included information about the dependencies that applications had on each other. For example, the Libraries' electronic resources reserve application (Ares) had been integrated into the university's course management system, Blackboard. That meant that when Blackboard was inaccessible, Ares was as well. All of these dependencies had to be taken into account when planning the schedule for the move.

Disadvantages to the Libraries Moving to the University Data Center

The Libraries have noted few disadvantages to moving to the new data center. What might seem like disadvantages are in reality just a change in the way we do our work. For example, we have been asked to inform someone in ITS before we go to the new data center to work on a server. This is a simple step, and has not hindered our work at any time. Another change is the need to use a tool created by ITS to configure our virtual servers, and we found that the tool has been configured to give us fewer administrative options than what ITS staff have. This has reinforced our understanding that we need to be present and proactive in representing the Libraries' interests in managing all of our computing equipment and software.

Migration Days

The majority of UAlbany's servers were moved from the main data center to the new one in ITB on August 9, 2014. However, we were unwilling to move all of the Libraries' servers on that day, which fell in the middle of the summer session. A compromise was reached between the Libraries and ITS that allowed many of the Libraries' less mission-critical servers to be moved on the same day as the university's servers. These servers were primarily ones that were used for development and backup purposes, one exception being the server that supported the Libraries' electronic reserves service. This server was dependent on the university-supported Blackboard server, which was being moved on August 9, so the Libraries' agreed to move this server that day so there would not be two downtimes for the electronic reserve system.

The Libraries' most critical servers were moved to ITB on August 18, 2014. This was the first day of intersession and would affect students and faculty the least. There were many people involved

in the move, including the Library Systems staff, the migration consulting firm staff, the professional moving company that was hired to carry out the move, and ITS staff who were responsible for the network and other support. Move activities included shutting down and backing up applications, powering off the servers, and packing the equipment. At ITB the equipment was unpacked, placed in its assigned rack location, plugged in, and powered on. Then each server had to be started, and applications tested. All of this activity began at 3:00 a.m. and continued until early afternoon. The day concluded with a conference call between all parties involved to confirm that everything was up and running as expected.

Lessons Learned

Participate in the Process

The Libraries were invited to participate in the planning for a new data center early in the process. ITS, UL, and other units with significant server collections met and discussed their computing needs and respective computing infrastructures. Once the construction of ITB began, the planning ramped up and monthly meetings of stakeholders became weekly meetings. Agendas for these meetings included round robin reports about

- construction project oversight;
- migration consulting;
- partnerships (with other units on campus, including the Libraries);
- status of our alternate data center (housed 10 miles away at another SUNY institution);
- campus fiber network;
- internal wiring and network design;
- administrative computing planning and move;
- research computing planning and move;
- systems management (storage and virtual environment) planning and move;
- data center advisory structure; and
- campus notification and public relations.

These meetings gave us an opportunity to learn about and understand all aspects of the data center migration project. Participants reviewed project timelines and other documents that were housed on a shared wiki space. After the data center migration consultants were hired, they began to use the Microsoft OneDrive collaboration space to share and distribute documents. Meeting regularly with all project participants allowed us to ask questions to clarify priorities and timelines and to advocate for the Libraries' needs.

Review Schedules Carefully

As with many construction projects, unexpected delays in the construction of the data center delayed all of our plans. Originally the building was to be completed in February; this was later

changed to April and then May. After the construction was complete, the building had to be commissioned, which means that every system within the building had to be tested independently by outside inspectors. Coordination of this work is very time consuming, and the completion of the commissioning delayed occupancy by another few months. The university was finally given permission to move equipment into the data center in July. In the meantime, our consultants were working feverishly to develop timelines for the move, identify, and secure a contract with a professional IT moving company, and create “playbooks” for each move.

The playbook is a document that includes

- names and contact information of everyone involved with the move;
- sequence of events: an hour-by-hour description of all activities;
- server overview: including the name, make, model, rack location and elevation, and contact person for each server; and
- schematics of both old and new server locations including details about each server as well as rack locations and elevations.

Library staff became concerned when the original date scheduled to move most of the university’s application servers, including the Libraries’ ILS server, was in the middle of the summer session. Although the projected downtime was only to be twelve hours (and probably fewer), library staff were not willing to have twelve hours’ downtime during a short four-week summer session. There were concerns that downtime, not only to the online catalog, but also to all of the Libraries’ databases, the website, online reference service, electronic reserves, and other resources would present a severe hardship to faculty and students. We also recognized the risk, however small, of something going wrong during the move that would cause a lengthier downtime.

At the same time the university was concerned about pushing the move too close to the start of the fall semester, as well as the increased cost of scheduling a second move date. During these negotiations it became apparent that the Libraries’ needs are different from administrative computing needs. Whereas the middle of a semester is a poor time for Libraries’ servers to experience downtime, it can be a better time for administrative computing, which is often busier during intersession when grading reports are being run and personnel databases are being updated. Ultimately, the Libraries advocated for and secured an agreement for a second move date, scheduled for the first work day after the end of the summer session.

Similarly, ITS was encouraging all of its partners across the campus to move as much computing as possible into their virtual environment. This is a worthwhile goal, but again the Libraries had to negotiate to make this change according to the schedule best for the Libraries and its users. The ITS virtual environment was a more current release of the Virtual Machine (VM) software than the Libraries were using, so the Libraries were faced with not only a migration, but also an upgrade. Ultimately, we postponed the VM migration until after the physical migration, and we have

benefited from waiting. Other partners have had to work through a number of kinks in the process, and the Libraries' VM migration has benefitted from the other partners' experience.

Clarify Costs of Centralization

When UAlbany began to consider and plan for a centralized data center, one of the concerns raised by the various data center managers from units other than ITS was the cost of centralizing their servers in another location. Centralized data centers have many costs: heating, cooling, security, staffing, cleaning, backup energy sources, networking costs, and more. The question on everyone's mind was who was going to pay for these costs. Would each unit have to pay toward the maintenance of the data center? Some objected to the idea of having to pay to be a tenant in a centralized data center, when they already had their own data center or server room at what seemed like no cost. The only cost they experienced was an opportunity cost of what else they could use the server room for. In the Libraries' case, the server room could be used for group study, office space, or other purposes, but it did not cost the Libraries money to use it as a server room because utilities are covered centrally by the university. On the other hand, by migrating some of our computing to the ITS virtual environment, we may save money in the long run because we will not have to replace hardware and pay warranty fees.

After much negotiation the university settled on a five-year commitment to no charges for the partnering units on campus, including the Libraries. This agreement was documented in a partnership agreement drafted by a group of representatives from all of the key units involved.

Contribute to the Development of a Service Level Agreement

Library staff contributed to the development of a service level agreement (SLA) for our participation in a centralized data center. Having an SLA in place ensures that all parties to the agreement understand their rights and responsibilities. We began by searching other universities' websites for samples of SLAs, which we shared with ITS staff who were assigned to this project. The establishment of a centralized data center includes several major elements: Data Center as a Service (DCaaS), Infrastructure as a Service (IaaS), as well as the network that connects it all. The SLA that was developed, still in draft form, has elements that address the following:

- The length of the agreement
- Network uptime
- Infrastructure as a service
 - Server/storage environment and technical support
 - Access to IaaS
 - File backup and retention
 - Maintenance of partner systems
 - ITS scheduled maintenance
 - Data ownership, security, responsibility, and integrity

-
- Business continuity, tiering, and disaster recovery
 - Availability and response time of ITS staff
 - Data center as a service
 - Environment and support
 - Building access and security
 - Physical rack space
 - Deliveries
 - Scheduled maintenance
 - Communications
 - Glossary

We recommend that institutions considering data center consolidation projects complete their SLA and other agreements before moving servers into a shared environment. In our case, however, we were unable to finalize the SLA prior to the actual move. This was not because of any particular demands that the Libraries were making, but was primarily because of the rapid approach of the deadline for moving into the new data center. It had to be completed before the beginning of fall semester, and preferably with a few weeks to spare in case anything went wrong. While the planning for the data center construction and migration seemed to stretch over a long period of time, the final few months turned into a frenzy of activity that ranged from last-minute construction details to nailing down the exact order in which thousands of pieces of equipment would be moved. Although not every detail was ironed out at the time of the move, the intentions and spirit of the SLA have been documented and it will be completed during 2015.

Communicate Developments and Plans

During the planning and development for the data center migration project, we recognized that it would be important to communicate any changes to the Libraries' systems availability to our users. ITS also recognized the need to communicate such changes. Both ITS and the Libraries took a many-pronged approach to communicating developments and plans related to the migration.

Within the Libraries we shared updates at Library Faculty meetings as well as meetings of the Library Policy Group (the dean's administrative policy team). We sought feedback from many groups on proposed move dates, establishing intersession as the preferred time to move any Libraries' servers that would affect access to resources used by faculty or students.

As the moves got closer the communication efforts were ramped up. Within the Libraries, we posted alerts on the Libraries' webpage that linked to charts indicating what services would be unavailable and when. We also included slides on the Libraries' main webpage with the same information. The same slide was posted on all three Libraries' flat-screen monitors, on which we post important news and dates. We sent mass emails to all Libraries' staff that reminded them when services would be down. Staff members who were responsible for specific services made an effort to contact their customers directly. For example, the head of Access Services contacted

faculty members about the scheduled interruptions of Ares, our electronic resources reserves system. Some of the downtime affected just users, and other downtime also affected staff who could not work in the ILS during the move. We planned alternate activities for staff members who could not work during the down time and had a productive Division Clean Up Day instead.

ITS also made great efforts to communicate to the university community about the moves and any potential downtime. Their efforts included mass emails to all faculty, staff, and students. ITS created and posted slides to the Libraries' flat screen monitors, as well as other monitors throughout the university. ITS also formed a team of liaisons from each school and college, using that group as yet another conduit to communicate changes. They shared draft schedules, seeking input on the effect of downtime on the university's functions.

Leverage Economies of Scale

One of the challenges of maintaining a distributed data center environment is that each system administrator or unit had to manage its own servers singlehandedly. In the case of the Libraries at UAlbany, we had moved in the direction of using the power of virtualization to manage many of our servers. Virtualization refers to the process of creating virtual servers within one physical server, thereby multiplying the value of a single server many times. The Libraries had virtualized a number of library servers, saving money by not having to purchase additional costly physical servers. However, ITS, with its greater purchasing power, was using more current and advanced virtualization software, hardware, and services than the Libraries. ITS created a suite of services that allows system administrators access to the virtual environment so they can manage their virtual servers from their own offices. By moving into their virtual environment (IaaS), the Libraries are able to leverage the economies of scale presented by their environment.

CONCLUSION

The consolidation of distributed data centers or server rooms on university campuses offers many advantages to their owners and administrators, but only minimal disadvantages. The University at Albany carried out a decade-long project to design and build a state-of-the-art data center. The Libraries participated in a two-year project to migrate their servers to the new data center. This included the hire of a data center migration consulting firm, the development of a migration plan and schedule for the physical move that took place late summer 2014. The authors have found that there are many advantages to consolidating data centers, including taking advantage of economies of scale, an improved physical environment, better backup services and security systems, and more. Lessons learned from this experience include the value of participating in the process, reviewing migration schedules carefully, clarifying the costs of consolidation, contributing to the development of an SLA, and communicating all plans and developments to the Libraries' customers, including faculty, staff, and students. As other university libraries consider the possibility of consolidating their data centers, the authors hope that this paper will provide some guidance to their efforts.

REFERENCES

1. "Fast Facts," University at Albany, accessed March 31, 2015, www.albany.edu/about/about_fastfacts.php.
2. "Data Center Consolidation, IT Consolidation," Webopedia, www.webopedia.com/TERM/D/data-center-consolidation-it-consolidation.html (accessed March 31, 2015).
3. Uptime Institute, accessed March 31, 2015, <https://uptimeinstitute.com/TierCertification/>.