Adventures in (Small) Datacenter Migration

Jon Kuroda, Jeff Anderson-Lee, Albert Goto, Scott McNally
University of California, Berkeley, EECS
{*jkuroda,jonah,goto,mcnally*}*@eecs.berkeley.edu*

Abstract

In May 2011, we embarked on an ambitious course – in 3 weeks: clear out a small, soon to be demolished, research datacenter containing 5 dozen research systems spanning 5 research groups and, along with a new faculty member's systems located off-site, move it all into another space suffering from 18 years of accumulated computer systems research history. We made it happen, but only after intensive pre-planning and after overcoming a number of challenges, both technical and non-technical, and suffering a moderate amount of bodily injury.

We present an account of our adventures and examine our work in facilities, networking, and project management and the challenges we encountered along the way, many of which were not primarily technical in nature, and evaluate our approaches, methods, and results to extract useful lessons so that others may learn from our reckless ambition.

Introduction

In late 2010, the EECS Department revisited previously shelved plans to renovate the north half of the 5th floor of the Computer Science building, Soda Hall, including demolition of 600 ft^2 of datacenter space ("530") shared by 5 active research groups and 1 defunct group. Mercifully, the new plans left alone a network closet that had been originally slated for a relocation but proved far too costly to move.[1] The demolition schedule shifted a bit but, by mid-March, eventually settled down around late May/early June with a construction start date of 2 June announced in late April.

Concurrently, we had a new faculty member bringing a rack full of research systems from a nearby Industrial Research Lab in Downtown Berkeley ("IRB") that needed datacenter space and needed to move by 25 May. We also folded this additional, smaller, migration into the overall plan.

With a larger campus and departmental re-examination of the utilization of datacenter space, we also saw this as an excellent opportunity to clean up, reorganize, and plan to upgrade our remaining datacenter space.

Others have examined the topic of large datacenter-scale change, most specifically Cha who looked primarily at the computing side of migration, especially system configuration, while we tightly controlled the amount of system configuration change and were also heavily involved in the facilities/physical plant side of the migration.[2] Similarly, Cumberland focused exclusively on system installation and configuration while our systems were already up and running and could not be wiped arbitrarily.[4]

Location Location Location

In evaluating new locations for migrating systems, we considered 5 major characteristics:

- Air Conditioning (measured in tons[2])
- Power (measured in kVA[3])
- Space (measured in number of racks)
- Existing Network Access
- Ease of Moving Systems

We considered 6 locations for evaluation, but only one made sense for relocation – a 1000 ft^2 datacenter space on the fourth floor of Soda Hall ("420A"). It had several points in its favor, including the most surplus cooling capacity and, after a number of upgrades, power, some pre-investment in overhead fiber distribution systems[4], sufficient connectivity to relevant networks, and physical proximity to related research groups, staff, and 530. Other facilities were either already or about to be filled to capacity, poorly suited for experimental systems requiring frequent physical access, accessible by too large a group of people, or lacked sufficient network access.

Like all such facilities in the building, both used raised floors with underfloor forced air cooling sharing space with underfloor power distribution and legacy network fiber runs. They were similar in most ways with 420A being a larger version of 530.

Table 1 summarizes the nominal capacity and actual usage of 530, IRB, and the final destination, 420A. At first glance, these numbers seem to indicate that moving these systems into 420A would run very close to if not right up against the rated capacity on Cooling and Power, but table does not take into consideration the later removal of defunct systems and the occasional but generous rounding up of usage numbers by Facilities staff.

Facility		AC tons	kVA	#Racks
530	Rated	30 (2× 15)	50	15
	Used	10	25	5
IRB	Used	2-3	10	1
420A	Rated	30 (2× 15)	100	34
	Used	13	45	12

Table 1: Facilities Utilitization Summary

While the best (or least-worst) choice, 420A still had several points against it, all centered upon its age, including an 18 year-old raised floor that had never been cleaned, 18 years of accumulated cable tangle stemming from minimal management of underfloor power and legacy fiber distribution, a multitude of circuit types instead of a single standard, 18 years of systems research history (aka "junk"), and a disorderly mix of standalone cage racks and relay racks.

Though the Computer Room Air Conditioners (CRACs), essentially in-room chilled water heat exchangers installed in pairs in each room, were nominally up to the task of handling the additional load, they too were 18 years old and not running at maximum efficiency. In fact, the manufacturer sold off that division shortly after Soda Hall opened in Fall 1994, and we are no longer able to get manufacturer replacement parts such as logic control boards – many of these units have had custom replacement boards installed by a third party vendor. A mix of "ownership" issues (the campus physical plant, not the department, manages the building HVAC system including the CRACs) and budgeting issues (the availability of funding for operational expenses versus funding for capital expenses in physical plant's budget) complicate outright replacement.

Power distribution in both locations consisted of an in-room PDU taking 3-phase input feeding under-floor runs of both rigid metal and armored flex-conduit (referred to as "cable snakes" locally) carrying single phase circuits ranging from 15A to 30A and 120VAC to 208VAC. After 18 years of use by a succession of resource-hungry computer systems research endeavours (each with its own power requirements) with minimal efforts to manage the power distribution, "cable snake" became an increasingly accurate term as the underfloor area devolved into an increasingly difficult to manage tangle of flexible conduit with many circuits disconnected but left under the floor, interfering with orderly air flow and creating a maintenance nightmare for the current staff.

We could address these problems given enough effort, which we had; time, which we had in short supply; and funding, which we had but is complicated in our academic environment though eventually tractable given enough time. For our immediate needs, this facility badly needed a thorough cleaning, a complete electrical survey before making any needed changes, and another thorough cleaning – these would be our immediate priority while other concerns would have to be dealt with as longer-term projects.

Timing is Everything

In our academic environment, we try to schedule major work involving downtime after May for many reasons. Finals, class projects, post-semester research retreats, conference/journal submission deadlines, and even VLSI tapeout schedules can all key off of the end of the semester, but we can engage in major work with relative impunity in the brief 2-3 months between the spring and fall semesters. Unfortunately, this also applies to major construction work, so we had to choose which of March, April or May would be the least disruptive time to accomplish this feat.

We went with May, partly due to circumstance, partly by design. Availability of the department electrician and a trio of work-study student staff would prove crucial, but they were committed to other work until May. Looking for an upside to this, we found this gave us more time to prepare and plan so that, when May rolled around, we could spend more time working instead of backing out of costly on-the-spot decisions made with little forethought or waiting to work because we had not thought something through extensively enough.

This choice had obvious downsides. While the demolition schedule carried a 2 June deadline, 21 May proved much more relevant due to off-site research retreats and long-scheduled staff travel in the last two weeks of May which gave us 3 weeks of time with the entire team present to do the actual work of prepping the new space and moving systems while power shutdown of 530 was scheduled for 25 May to allow for dismantling and disposal of the two CRACs. This aggressive schedule came back to bite us once or twice, but the hard and very real deadline proved to be very strong motivation for us and gave us greater ability to cut through bureaucracy –

pushback from those outside our team or claims that we "could just get an extension" were rebuffed with a reminder that contractors were arriving on 2 June and that delay would hold up a major construction project, and, if necessary, an invitation to discuss the matter with the department chair. This, however, only happened once with a colleague who was unaware of the entire scope of the work and was quickly handled.

Our faculty, in particular, left us alone to do this work and did not question the migration schedule. While they already trust us in general on operational matters, they had specifically been informed of this work by the department chair and the facilities director to prevent exactly these sorts of questions – there was faculty buy-in to this adventure before it even became our concern. While we may not have particularly wanted to go down this road, at least this road had already been paved for us. Of the 2 other active groups who had systems in 530, one group would be moving into the renovated space and the other had only a small handful of systems and did not mind much as long as the systems were back up eventually – we did not anticipate nor receive any pushback from these two groups on the migration schedule. That we received no pushback on schedule from faculty or users and almost none from staff still amazes us.

People Get Work Done

One immediate challenge we faced was that we have no staff dedicated to Datacenter Management or who have it as a primary job responsibility. Instead, we have a number of staff who do work in datacenters as one aspect of their jobs, whether they be systems administrators, facilities managers, or network administrators. Most notably, the systems administrators responsible for completion of this project all work directly for 3 of the 5 research groups affected by this move – the other 2 research groups did not have systems support staff to contribute to the overall migration.

Our team consisted of one Department Facilities Director, one Department Electrician, three systems administrators with deep institutional knowledge, one new guy who started in May, and three work-study student helpers. Missing from our merry band was a network administrator to handle the myriad of network changes needed for all this – primarily a number of changes in VLAN assignments. The staff member most familiar with the network involved in this migration had taken another position elsewhere on campus in February 2011 and left behind another network administrator unfamiliar with both the overall topology as well as the platform specifics. This hole would come to haunt us later and nearly derailed the migration schedule.

Close ties developed over the past decade between team members proved vital to success given our tight timeline – having to go "through channels" for change requests and having to continually re-establish a shared terminology would have crippled our ability to move quickly on a tight schedule. We did, however, observe a marked discrepancy in the preferred manner of communication – while we all relied on 1-to-1 in-person communication for low-latency high-bandwidth communication, we never converged on a single mailing list, wiki, or any other particular form of asynchronous collaborative "hivemind". Periodic synchronous standup meetings proved to be only the consistent way to keep us all on the same page.

Goals

With limited staff and time, we had to keep our goals modest while at the same time ensure that we allowed for future facility improvements and upgrades. As noted, we had already decided not to engage in major upgrade or reorganization work in 420A. We also had to decide how much support to give to systems that did not belong to our faculty but instead belonged to the two groups without systems administrators.

In the end, we settled on providing a minimum baseline level of service of rackspace, power, and networking for all systems migrating out of 530 but only systems that belonged to our faculty received hands-on support from us. We convinced the Department IT Director to take responsibility for systems that belonged to a defunct research center whose sole faculty member had retired years earlier. Of the two groups lacking systems administration staff, one chose to hire the department's User Support Group to handle the hands-on migration work while the other gave the work to one of their undergraduate interns.

For our own faculty's systems, we settled on 3 service guarantees:

- Max of 1 downtime/system
- Max of 1 day/downtime
- Minimize impact on deadlines
  - do not move everything at once

We aimed to have systems back up within 24 hours after taking them down for migration and, once we said a system was back up, for it to stay up barring user needs or "normal" routine operational needs such as periodic OS patching. We particularly wished to avoid taking systems down again to move after announcing that a system had been moved and was back up.

The last goal proved to be the most complex but also the one most beneficial to us. We decided early on that moving everything all at once in one fell swoop was far too disruptive to our users and risky for us if we took a wrong step, so we instead chose to move systems based on when relevant user populations needed them – moves more than 2 weeks before a deadline or the day after a deadline were acceptable, but not during the two weeks before a deadline. This would ultimately benefit us as we could pick and choose unused and therefore less critical systems to move first in order to test the waters after which we could move systems in larger groups.

For the new faculty member's IRB systems, we focused on:

- space for racked systems
- installation of electrical circuits
- transport on or before 25 May

We were not immediately concerned with getting the systems from IRB up and running, only with making sure that there was a location for them and making use of the electrician's time while we still had access to him in May.

When possible, we allowed our decisions to be guided by the pursuit of progress towards a cleaner, more well organized datacenter space, but were prepared to make well-defined and easily undone short-term decisions in order to meet our 21 May deadline.

Space Planning

We had actually begun planning for this over a year prior when the renovation plans first came across our desks. Already dissatisfied with the collection of ad hoc changes made to datacenter spaces throughout the building and the way that they hampered any growth or reorganization, we all saw this as an opportunity to rip out as much cruft, junk, and accumulated history as we could manage in whatever time frame we could acquire.

While we could not muster enough momentum or staff time to accomplish significant datacenter cleanup after the department shelved these initial plans, the ideas for cleanup and upgrade were still fresh in our minds and on paper when the department took the renovation plans back off the shelf. Additionally, we had already done a survey of 530 and 420A in late-2009 as part of a campus datacenter utilization survey, so we had a good handle on who had what systems in 530, how much power they consumed, and how much rackspace they needed.

Initial migration planning started with a overall survey of 530 to review any major changes since the 2009 survey. We paid special attention to the type of circuits we would need in 420A – while individual systems used standard IEC 60320 electrical connectors such as found on typical PC systems, our in-rack power distribution used a variety of means to connect to building power including 4 different NEMA twist-lock connectors. We conducted a similar general survey of 420A to confirm general impressions of the space and to note things that we could correct before May without the assistance of the electrician or the need for the trio of work-study students.

We briefly entertained the notion of installing an overhead busbar power rail system, as is now increasingly common in new facilities on campus, but quickly placed it on the "needs time and money" list. Within the time constraints we had, particularly the electrician's availability, we had to make do with more limited incremental changes to the existing underfloor power distribution system. Overhead power would become one of many recommendations that we would make for future datacenter upgrades.

We also held off on any major upgrade work to the air conditioning system, again for reasons of time. In place of capacity upgrades, we pursued two alternatives. First, the Facilities Director ran two day long experiments running 420A with 1 out of 2 CRACs shutdown to see if each of the 18 year old pieces of equipment could handle the existing load alone – which they did without failure. While not a strictly rigorous experiment, it demonstrated that we could have enough cooling capacity given prudent placement of systems and pruning of unused or offline systems. It also enabled us to pursue a second avenue – maintenance service and overhaul. The Facilities Director scheduled two maintenance periods for each CRAC involving aggressively proactive replacement of worn parts, cleaning of water piping to and from the building's rooftop chilled water supply, and servicing of each CRAC's 3 compressors. Our Facilities Director estimates that this restored about 20-30% of efficiency back to the CRACs though he notes that building AC makes exact numbers difficult to obtain – in warm months, building AC runs more often, creating a shell of cooler rooms surrounding 420A while in cooler months overall need for heating is rare due to the effectiveness of the building's own insulation.

We reviewed several ways to rearrange 420A, but ultimately retained the existing arrangement of 4 main rows of racks and 1 catch-all row with a more concerted effort at enforcing hot and cold aisle separation for dense installations while relegating less dense installations to "warm" aisles. We did rearrange rack allocations so that projects, which previously had equipment strewn across various disparate racks due to the floorspace equivalent of disk fragmentation, could benefit from physical proximity, thus alleviating network fiber distribution complexity as well as strengthening project and group identity. In addition, we identified equipment from prior projects which had been abandoned in place and was eligible for reuse or salvage.

Once we established an initial rack-by-rack layout, mapping research groups to racks, we worked out which systems would go into which rack and ultimately came up with a rack-unit (RU) by rack-unit layout. While we could not move any systems until May and knew that plans could change in an instant, this gave us a start on planning in-rack power and network wiring and allowed us to organize systems in a more sensible fashion compared to the previous “which rack has room?” method.

After running through a few different ways of sorting systems into racks, we eventually identified 3 classes of systems that aligned naturally with their owners and purposes that also lent itself to a means to organize the racks and to answer the inevitable question of ”where do I put this system?”

Experimental research systems

Our faculty’s systems
Sat on “Research” network
Novel Hardware is the point
Often had two of each kind

“Production”ish research systems

Our faculty’s systems
Sat on “Research” network
Stable research platforms/services
Clusters, OS dev, storage

“EECS” systems

(mostly) Other faculty’s systems
Sat on “Department” Network
Group web servers, SW

Our eventual rack layout would later reflect these alignments and would let us make use of some limited “luxury” hardware resources more effectively. For instance, while we did not have the resources on hand to put a UPS in every rack, we did have a new-in-box UPS that had been bought but was never used for a now-decommissioned system – we installed that UPS into the Production rack to provide battery backup to a 24TB storage system and to some management systems. Similarly, other racks got “intelligent” power strips with remote outlet control and per-outlet power metering that, while not so useful for research quality results, let us keep tabs on spikes in power consumption as researchers ran experiments. Meanwhile, racks housing systems that we expected would see frequent hardware reconfiguration got an in-rack KVM so we could avoid having to un-rack systems just to get a console on them. Most of this hardware (cabinets, power-strips, UPSs, KVMs) was re-purposed from prior projects.

As the end of April approached, we began more definitive preparations. One admittedly sneakier one was to make daily sweeps of systems in 530 to check for running user processes and when the last non-staff login occured. If no user processes were running and the last non-staff login was more than a few weeks prior, we preemptively shut the system down. We shutdown 15 systems this way and only had to turn one back on – the remaining 14 remained powered down until we moved them at our relative leisure in May.

Some of the more obvious space preparations included basic cleanup of the space – we lost track of the number of cardboard boxes we had found squirreled away in every possible corner, nook, and cranny – collection and removal of abandoned or deprecated systems that had been left behind in racks but never tagged as excess, and tagging of equipment for storage and later repurposing. In preparation for the inevitable exodus of unused systems from both 530 and 420A that nobody was quite prepared to send to the campus “Excess and Salvage” unit, the Facilities Director had begun his own cleanup of a large basement storage room for our use during the move.

By mid-April, we had a good handle on the work needed to ready 420A – outside of a fair amount of hands-on physical labor, we saw no major facilities obstacles to having space ready for move-in during May, leaving only networking left as a major concern.

Network Planning

We had three areas of concerns regarding the network changes needed to support this move: the changes needed, who would do the work, and, as usual, the time available. As with planning for other parts of this move, we chose to stick with minimal changes instead of an overly ambitious redesign.

The systems moving out of 530 were spread across two distinct networks, a ”Department” network, managed and funded centrally by the department, that provided general commodity network connectivity throughout the department and a “Research” network, funded by research grants and donations and managed by research systems support staff until early 2011, which evolved out of a wider network deployed for a campus-wide clustered computing project to serve the needs of specialized computer systems research in the Department. The vast majority of systems belonging to our faculty sat on the Research Network while the dozen or so systems that belonged to other faculty sat on the Department network.

This presented one small but immediate problem. While 420A historically supported systems associated with clustered and distributed computing research – the projects involved in such research provided their own network hardware to support the higher density networking they required – the Department network had very little presence in 420A at all, no more than a dozen ports available via network drops pulled in from a nearby network closet that were meant for one-off systems, not for higher density installations. Department networking staff did not have any spare equipment to install a managed switch to support denser installation in 420A of systems on the Department network, but fortunately research networking had enough spare equipment to loan out a switch which someone would setup as a managed switch attached to the Department network. The last detail about it being used as a managed switch would change, but this plan would remain otherwise unchanged.

A larger question was what direction to take the Research Network's presence in 420A. The Research Network's presence in 420A, once extensive to support large clustered computing projects, had itself dwindled in size as systems' power and cooling density rose far faster than their space and networking density[5] and by this time had evolved into a few network stubs supporting smaller projects.

One such stub, a group of 4 daisy-chained switches attached to the Research Network via a lone 10Gb/s link over long-range fiber, was on the right VLANs which opened up the possibility of just daisy-chaining even more switches. We had previously discussed plans to stem the growth of the Daisy Chain of Doom (DCOD) but lacking sufficient spare long-range optical modules for a second long-range fiber run, extending the DCOD was straightforward, predictable, and cheap since we did have plenty of switches and short-range optical modules. We understood the downsides of relying on what would turn out to be a daisy-chain of 7 switches, but felt that it would be acceptable for the short-term (6-9 months) until we could spend enough time planning more extensive changes.

Regardless of any physical topology changes, we quickly realized that there would be a significant number of changes to port VLAN assignments to support systems with private interfaces on a separate VLAN. This led to the biggest question – who was going to do the work of reconfiguring the switches?

Prior to 2011, research systems support staff shared management access and duties with a lead "Network Guy" who himself had worked in our team as a split network and systems administrator before transitioning in 2009 to a full-time network mangement position supporting both the Department and Research Networks. Upon his departure, he handed over the Research network to the remaining network administrator at the direction of the department IT Director who wanted to see both networks managed in a more unified manner.

We were wary of this change, specifically losing access to manage the research network, but, in a good faith attempt to support the IT Director's direction in network management, we went along with his request that we send all of our network change requests to the department network staff which consisted of the remaining network administrator and the soon-to-retire infrastructure services manager. While we expected some communication and culture differences, we did not anticipate at this point the delays that would come with this and nearly derail the migration.

In late April, we met with the remaining network administrator to go over high level plans and examples of the changes we wanted so that everyone was on the same page. We held one more meeting to go into further detail and left with everyone understanding what we needed and when we needed it. At this point, every major task was identified and assigned to one or more persons.

Progress

Once May started, we gained access to the department electrician, work-study student labor, and, much to our delight, a large storage room courtesy of the Facilities Director for anything and everything we wanted to remove from 420A or 530 but were not quite ready to junk. Also joining us on 2 May was our new systems administrator who showed up for work right as we began the bulk of the work. We all met on 1 May for an initial meeting to confirm that we were all on the same page, and from that point, work progressed quickly.

Our first order of business was a detailed electrical survey of 420A so we actually had some idea of what circuits were actually live. At the same time, we were tagging equipment either to go to Excess and Salvage or to storage for the work-study crew to remove from 420A which they did as fast as we could tag it. Within a week, the electrician had mapped out all circuits in 420A, made all of the initial electrical changes that we had requested, and disconnected all hardwired powerstrips that had been attached to the legacy relay racks. We would later ask for a few changes which were quickly handled. Once this work was completed, the work-study crew set to work removing relay racks so we could bring in cage racks from storage that we had setup in a staging area with in-rack power distribution and cable management.

In the second week of May, we had begun to bring in networking to individual racks. While we waited for the network administrator to configure switches to add to the DCOD attached to the Research Network, we worked on bringing in more access to the Department Network by installing a spare L2/L3/L4 switch in the rack we had setup for the two groups whose systems all sat on the Department Network.

As we had anticipated, setting up a optical fiber link back to the Department Network did not work out – we only had 10 gigabit optical modules for our equipment and the Department network staff did not have spare 10 gigabit optical modules to support a fiber link to our switch, only 1 gigabit modules. Instead, we located a free copper network jack in 420A, connected the switch to that, and proceeded to setup a pricey[6] L2/L3/L4 switch as the functional equivalent of a simple 48-port L2-only desktop switch. Though arguably a waste of an expensive piece of network hardware, this temporary loan allowed the department networking staff the flexibility to later provision a switch from their current vendor, get the proper optical modules, and configure it more fully to allow them to bring in multiple VLANs as needed.

At the end of the second week of May, we were ready to let the two groups move their systems into 420A. We wrote up instructions on how to get their systems up and running in 420A and distributed this the next Monday – they would be the first people to migrate out of 530. We would not be able to move our faculty's systems out of 530 until the third week of May due to delays in getting our network changes handled.

Network Lag

In the second week of May, we experienced slower progress with our network requests than we had expected. Though we estimated the actual work would only take an hour or two, we anticipated some delay due to forseeable factors on the part of the network administrator such as lack of familiarity with the Research Network, a strong desire to get everything setup exactly right for us, and an already busy work schedule. However, it became increasingly clear that the delay would be well beyond what we anticipated or could handle given the short timeframe.

Our first experienced a short delay when the network administrator pushed back a day on handling our network changes; though unexpected, we attributed it to a heavy workload due to taking on all day-to-day support for the Department Network after Network Guy's departure and felt that we could absorb this delay as we still had some facilities work left to finish.

Though the network administrator had worked with and trained on equipment from the vendor used by the Research Network, that experience had gone unused for a number of years due to the department's use of a different vendor. To help overcome this, one of our team wrote the switch configurations for the network administrator to load onto the switches, saving the network administrator time and transferring responsibility to us if something went wrong.

The final and unmistakable sign came near the end of the second week of May when we received notice that the reconfiguration of our switches, essentially loading the switch configurations we had written ourselves, had been reassigned to another member of staff who also had little to no recent experience with the network environment. It was at this point that we realized that something much more fundamental was going on here than just lack of time on the part of the network administrator. At this point, we only had a week left with all team members present to complete the work – any further delays would irreparably derail the migration.

A hour-long meeting with the IT Director revealed that he had been unaware of our previous access to manage the Research Network and of our relevant expertise and that he had also partly based his decisions on inaccurate information about relevancy and recency of other staff members' experience. We appreciated his desire to see both the Department and Research networks managed as a more unified single entity and understood his concerns that restoring our access to manage the Research Network could lead to further divergence, but we could no longer tolerate holding up the schedule to accomodate this nor wait for staff who weren't familiar with the systems involved. We reminded him of the lesson from from Brook's "The Mythical Man-Month"[1], namely that adding more staff to a project running behind schedule [or, as we added, running very close to it], particular staff unfamiliar with the project, will cause it to fall further behind schedule.

By the end of the meeting, we had, albeit with caveats about keeping the department network administrator informed, regained administrative access to the Research Network hardware and, with passwords literally in hand, we walked back up to 420A and started reconfiguring switches. That evening, we were up and running and ready to start moving our faculty's systems into 420A that night – on-schedule completion once again appeared within reach. In the weeks and months after the migration, we would return to the question of why it took so long to achieve our network goals, regardless of how we achieved them.

Back to Work

With our network problems resolved, we returned to the exodus from 530. With a few days of work time lost to dealing with the network management problems, we had to toss out a fair bit of our move schedule and start doing a lot more ad-hoc scheduling. Luckily, when we first started working out downtime schedules with users, instead of just setting specific dates, we also asked for "OK" windows, akin to space launch windows, during which it would be acceptable, though maybe not optimal, to shut a system down with short notice. This allowed us to aggressively pursue a more dynamic schedule where we check systems for user processes, quickly check-in with users about short downtimes that morning or afternoon, and, barring any objections, move systems in as large a batch as we could manage with downtimes hovering around an hour for each system. At times, users would even proactively inform us that they were done with a system and could bear to be without use for some period of time, thus saving us the work of regular polling.

By the middle of the third week, we were back on track and even a little ahead of schedule as systems were being ferried down from 530 to 420A 3-4 times a day. By using a few spare and borrowed rack cabinets in 420A, we were able to more easily stage the migration of systems out of 530. The alternative of moving whole racks at a time (versus opportunistic migration of systems) would have been more challenging and would have allowed less opportunity for the reorganization of systems in 420A.

At the same time, we had managed to carve out time for one team member and the electrician to make a field trip down to Downtown Berkeley to confirm details of power and transport of the new faculty member's systems from IRB to EECS. By the end of the third week, we were down to just a few systems and lots of supplies and equipment left in 530 that were migrated to 420A or sent to storage. That weekend, one team member went out of town to handle site prep for an off-site research retreat followed by long-scheduled vacation and would not return to campus until two weeks later. While the IRB systems arrived on 25 May, we consider the day of completion for the migration to be 24 May which coincided with the first day of work for a newly hired computing Infrastructure Services Manager.

Lessons Learned

We learned a lot from this adventure. The one thing agreed upon by everyone involved as well as by several outside observers is a general sentiment of "Never. Again. Ever." This was probably near the limit of what could be accomplished with the resources we had available – and we came very close to failing. While we had options available for every problem we ran into, in some cases, some of those options had worse long-term consequences than failure. We got by on our own resourcefulness, persistence, and sheer luck to make up for a lack of room within which to fail. We hope that, by showing what happens when one runs a migration with the bare minimum time and staffing, others may take heed and step back from the proverbial cliff's edge that we danced upon for a period of a month.

A key overall lesson we kept coming back to was "Plans are just that – plans." The more time one has, the more one can try to bend reality to a plan – the less time one has, the more one must bend plans to fit reality. We lost track of the number of times some trivial problem arose that did not fit into our plan – like a supposedly 42RU high rack turning out to be 41.75RU high – instead of losing sleep over it, we found ways to adapt, like we did with the shorter-than-advertised rack by moving a server to another rack despite it not completely lining up with how we organized systems into racks. By looking at our plans more as initial guidelines or a starting point, we gave ourselves the freedom to deal with small unforseen problems without getting hung up about them and to look at larger problems not as problems but as course corrections.

The only truly disruptive problem was the delay in getting network changes made. While the problems we faced are getting a more formal review now that the new Infrastructure Services Manager is settled into his job, it is clear we had problems with communication, awareness of staff skillsets, and what could best be described as differences in culture. Others in the department have cited lack of project management, formal or otherwise, as a skill not explicitly present in our team – though our Facilities Director served as de facto project manager to keep track of progress and did communicate our problems to the IT Director, we found that the IT director did not understand the full situation until we had our own sit-down meeting with him. More explicit project management would likely have caught and handled these problems earlier. For our part, it is fair to say we could have been more direct and explicit about when we needed this work done and could have proactively listed restoration of our administrative access to the Research Network hardware as an alternative if a deadline were missed.

The question of "What would you have done if you had not received access?" has come up. We find it hard to believe that the IT Director would have said, "No." after being presented with the facts, but from a technical standpoint, it would have been easy – with physical access to the hardware, we could have the needed access via brute force methods.[3] This would have required disruptive power-cycling of each switch, but we could then do what was needed. From a management standpoint, it would have been contentious at best. At the very least, we would have informed our faculty and the Facilities Director of the problem and what we planned to do. It likely would also have perpetuated an image of us, whether rightly so or not, as "wild cannons" with little respect for authority and process who were difficult to work with and ultimately would have contributed to antagonistic relationships with our colleagues and our management. This was not something any of us would have considered casually.

Though not disruptive to the migration, still unfortunate were the moderate injuries suffered by staff. One team member, while working alone after hours with the raised floor, bashed his knee on two occasions, resulting in a bruised kneecap that took two months to heal completely. Had he suffered more serious injury, it was possible no one would have noticed until the next morning. His time working late would have been better spent planning and preparing for the next day's work or even sleeping. Another staff member from Shipping and Receiving broke two fingers while moving into a elevator a cage rack that was still loaded onto a pallet. Though trained in use of pallet loaders, he was unaware that racks shipped on pallets often come with ramps to aid in unload the rack to avoid having to move the entire package like that – while most systems administrators in the department have had to deal with this and were aware of this feature, we wonder how widespread this knowledge is among facilities staff outside of those who work extensively in datacenters.

While attributable to the rapid pace and aggressive schedule of the migration, these injuries were avoidable indicating that at times some team members pushed themselves too hard or that they should have taken a more active role in physical tasks delegated to other staff to share our knowledge regarding better methods.

Our team took to heart the first lesson from "The Mythical Man-Month" about bringing on new team members to quick-paced or behind-schedule projects. All members had known each for at least a few years, had worked together on other projects, and were familiar with each other's habits as well as with the institutional knowledge of how things worked in the department. We did not bring our new systems administrator onto the migration project until after the majority of the facilities work had been finished and we could give him tasks that could be done jointly with someone else. We consider the case of assigning network configuration tasks to staff unfamiliar with the environment to be an example of the "Mythical Man-Month" fallacy that Brooks describes.

Similarly, we confirmed what might be considered an obvious notion – that fast-paced schedules are no place for in-depth staff training. When the schedule is fast, and the room for slack is tight, this is no place to be bringing in people new to the environment. The training overhead of new staff is unaffordably high in addition to the extra communication overhead of adding an extra person, experienced or not. As with our new guy, we instead followed another lesson from Brooks – give more routine tasks to new staff to allow more experienced staff to tackle the more difficult problems.

One thing that surprised us but really should have been obvious was the lack of convergence on any asychronous collabarative systems. Use of e-mail lists were sporadic at best, and only technical staff made any use of wikis or other web-based systems. As noted, for overall team synchronization, the only consistently successful method was a periodic standup meeting. Even among technical staff, paper notes left on racks often served better for passing short must-read notices about a rack or a system while things were in flux while the use of a wiki appeared more well suited to noting the steady state once a system or rack had been successful moved into place and deemed stable. This is extension of our previous lesson – once a project starts, training project members on new systems is inadvisable at best unless the benefit is so large compared to the training overhead that failure to adopt the new system could result in project failure.

The lessons we learned could best be summed up as "More time, better communication, and slow down." The ambitious plan to completely clean up the target room (420A) was cut short due to lack of time. While the result was a vast improvement, the time required to perform the cleanup was underestimated. In addition, it was unclear at the outset what parts (shelves, bolts, power strips, cables, etc.) might be needed for the final configuration. An earlier start on the clean-up would have helped; at least more of the parts could have been sorted in advance for subsequent disposition.

Future Work

While we were able to accomplish our basic goals of moving all systems from 530 to 420A in the time available, we still have a great deal of work left to bring 420A and other similar facilities in the department up to more contemporary standards. Some of this work is tractable in the near-term while other work will remain the focus of longer-term efforts involving questions of funding and staffing.

Our short-term work will focus on continual efforts at clean-up in all facilities. The one constant we have found about any datacenter facility, especially in an academic research environment, is the tendency for old equipment to accumulate and linger around so long that people forget what a piece of equipment is for and, as a result, become afraid to get rid of it. The one constant we have found about buildings with both raised floors and drop ceilings is the tendency for those areas to become absolutely filthy, especially raised floor plenums which cause systems to take in large amounts of dust and "biological debris". Both situations require both immediate and ongoing attention to mitigation. We are currently evaluating in the short-term a number of different options for front-of-rack filters.

Mid-term work focuses on networking, in particular the hack that is the 7-switch-long DCOD currently feeding almost all of the research systems in 420A. Current options include a 16-port 10GbE distribution/aggregation switch or setup of a double-ended string of switches. For now, the DCOD suffices, but the physical path of the fiber is tortuous at best due to the ad hoc manner in which the original stub network in 420A grew with fiber criss-crossing the room multiple times. Additional switches will only complicate this further and limit growth. Related to that is the replacement of the switch currently deployed in a 'production' role on the Department Network. It should at the least be managed as a fully configured switch instead of setup as a $6,000 dumb L2 desktop switch, but ideally would be a model from the vendor currently used for the rest of the Department Network so that it can be more easily managed by department networking staff. Finally, we are very interested in finding ways to promote a more unified approach to network management that avoids the maintenance of two completely separate network domains.

Longer term work on the order of a year or more include several facilities upgrades. The first and foremost, but likely to take the longest is the replacement of the in-room CRACs which are nearly 20 years old and no longer supported by their manufacturer. Current measurements indicate that we are already using approximately 80% of our AC capacity while we still suffer a compressor failure about every 2-3 months in 420A. Replacement requres negotiation with our campus Physical Plant and could take a year if not more, but is neccessary to support future growth.

Other work we would like to pursue all relate to power distribution. The variety in types of circuits installed complicated work – standardizing on a single type of circuit, say 208VAC at 30A, would greatly ease management of the underfloor power distribution system and would make it easier to standardize on a single type or vendor of in-rack power distribution units (PDU). It is already hard enough to find PDUs with certain features – per outlet power monitoring, metering, and control, reasonably secure remote access, and a usable API for developing our own applications – that trying to account for even a handful of circuit types makes this impossible given current vendor offerings and impede efforts to gain better insight into power usage. We could pursue this work on a piecemeal basis, but we wonder if that would result in less standardization. Potentially the most extensive work we would like to pursue is transition to an overhead power distribution system. Though obviously costly, this would yield huge benefits in restoring orderly airflow to the underfloor air circulation space, easing of time costs of dealing with numerous circuits types and simplifying power usage surveys and audits.

On the non-technical side, we look forward to addressing clear deficits in key areas such as datacenter, network, and project management, communication and culture barriers between research and production operations, and management awareness of staff expertise along with staff awareness of management plans. We expect that this work will be ongoing for the rest of our professional careers – not because we think that we will always have these deficits but because the only way to avoid developing these sorts of deficits is by continually working to ensure that they do not develop.

Conclusions

We pulled it off, but just barely. We needed every single last day available in order to complete the work neccessary and could have used an extra day or two for breathing room. We got by on large measures of determination and dedication, resourcefulness, and sheer dumb luck. We look at this work as an accomplishment worth being proud of but also as an example illustrating all the things that one should have – many of which we did not – in order to embark on a similar datacenter migration adventure with a more reasonable chance of success and better options in case of something less than 100% success.

Acknowledgments

We would like to thank all the EECS staff who provided invaluable help without which this migration would not have happened, the faculty and researchers for their understanding of the pressures bearing down upon us and, of course, for all the caffeine that they made sure we had in ample supply.

References

[1] BROOKS, F. *The Mythical Man-Month*. Addison-Wesley, Reading, Massachusetts, 1975.

[2] CHA, L., MOTTA, C., , S., B., AND AGARWAL, M. What to do when the lease expires: A moving experience. In *Proceedings of LISA 1998* (1998).

[3] CISCO. Password recovery procedures. `http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_tech_no%te09186a00801746e6.shtml`, Jan 2007.

[4] CUMBERLAND, D., HERBAN, R., IRVINE, R., SHUEY, M., AND LUISIER, M. Rapid parallel systems deployment: Techniques for overnight clustering. In *Proceedings of LISA 2008* (2008).

Notes

[1] Rough estimates for network closet relocation ran in $500K range.

[2] 1 ton of cooling is 12,000 BTU/hr or 3,517 W. The refrigeration and air conditioning fields use this unit to denote the heat required to melt 1 short ton, 2000 lbs, of ice at 0 °C in one day, representing the cooling provided by daily delivery of 1 ton of ice.

[3] 1 kVA is 1 kilovolt-ampere and is used to measure "Apparent Power", the product of root-mean-square voltage and current. "Real Power", measured in watts (W), refers to the power actually usable by devices.

[4] This "system" amounts to a grid of overhead Panduit-style fiber trays that meets up with an in-room fiber termination box fed by a nearby network closet.

[5] This trend would peak in the mid-2000s with the installation of a 128-node Itanium2 cluster.

[6] The switch was part of a large donation, so we do not know how much it cost the vendor to give it to us, but the current street price with optics is around $6000.