

Duke Office of Information Technology

Minutes

September 15, 2005

Members present : Ed Anapol, Owen Astrachan, Pakis Bessias, John Board, Shailesh Chandrasekharan, Ed Gomes, Dick Danner represented by Ken Hirsch, Shiva Das, Tracy Futhey, Michael Gettes, Daron Gunn, Craig Henriquez, Billy Herndon, Deborah Jakubs, David Jamieson-Drake represented by Bob Newlin, Roger Loyd, Gregory McCarthy, Kyle Johnson, George Oberlander, Lynne O'Brien, Mike Pickett, Tong Ren, Rafael Rodriguez, Dalene Stangl, Molly Tamarkin, Christopher Timmins, Trey Turner III, Robert Wolpert, Steve Woody

Guests: Kevin Davis, OIT; Brett Walters, Franklin Center; Jason Doty, Franklin Center; Mark Olson, Franklin Center; Ginny Cake, OIT

Start time : 4:02 p.m.

I. Review of Minutes and Announcements:

- ? Julian Lombardi hired as AVP-ASTS: Tracy says I'm pleased to announce Julian Lombardi will be joining OIT November 1. He will be the Assistant VP for Academic Services and Technology Support. He will have responsibility for all customer-facing interactions such as the Help Desk and will be interfacing very closely with Lynne O'Brien and the efforts CIT does as well as people like John Harer in the Center for Computations Science, Engineering and Medicine . We look forward to having him on board.
- ? Lynne O'Brien says the Podcasting Symposium will be next week. Casey Alt and Richard Lucic have been the coordinators and have put together a 2-day symposium on various issues related to podcasting such as business models, copyrighting, and its cultural significance. We have had a phenomenal response to this, we never really advertised because so many people signed up. There are over 300 people coming from Australia to Canada . There is a website up about it at <http://isis.duke.edu/events/podcasting/> Mark Olson says registration for the symposium is currently closed, though it may reopen. We intend to webcast it.

II. Impact of articles/books online – our culture of printing - Deborah Jakubs, Chris Chin, Kevin Davis

Deborah says I'm introducing this topic by having Chris give the student's perspective. He's a DSG bigwig, and Chris has come to the library to talk to an executive group to express some concerns about printing and printing abuse. From the library perspective there are two issues: there has been a really dramatic increase in the amount of money we have been spending on paper. Replacing paper in machines and toner cartridges

are part of that. We are concerned about how we control that, how we regulate it. The ePrint process has helped that to some extent, and Ed Gomes has been very involved in that as well.

The immediate thing that put this on the agenda is the printing of e-reserves, which faculty show is a big help to students. The trigger was a 300-page textbook that hadn't quite been published yet that a professor put on e-reserve for a course. I first ran into this walking into the library and saw a line of students at a printer. I inquired, and they were waiting for the printing job to finish over and over again. We also have the faculty perspective, so how do we communicate to them what is appropriate to have on e-reserve? Expecting students to read online is unrealistic: it would be a losing battle to try to change that.

Chris Chin says in terms of years it becomes necessary to analyze different mediums through which information gets from a professor's syllabus to students. There seems to be combination of fair use and a desire for everyone at the university to be on the same page in same way: some think it is desirable, some not so much. Ways I've come up with: buying books; teachers doing their own photocopying, implying departmental resources; blackboard, e-reserve, e-journals, which are something libraries and others are encouraging; books on reserve, which these days more relegated to out of print materials, and is somewhat archaic considering most classes have more than ten people; the last is using library resources. It's important to look at this list and think of who is pushing what. Deborah asked me to discuss why students don't read things on the computer. You can't highlight, can't take notes, and can't bring it to class. Everyone in every department is using a different combination of these seven methods.

John Board asks what has happened to course packs?

Deborah says part of the reason those have gone away is because it was up to faculty to figure out the copyright for all the materials in the course pack.

Kevin Davis says on the students' part they have a very open printing system available, and they don't worry about a cost. It's a tough question to entirely answer how we get to this culture. Some speculation is that much of it is related to students printing flyers for parties and campus events – we've heard about this from some students. Also, it may be course-related. It's worth testing this out with last year's top printers, who printed about 7,500 pages each. ePrint has reduced waste, but there has still been a growth in printing. There has been a 50% increase in printing in the year after ePrint started.

George Oberlander asks are they heavy printers because they print many medium size jobs or because they are printing large jobs? Do we need a mechanism to limit the size of the job?

Lynne O'Brien says two factors I haven't heard mentioned that motivate putting materials online are flexibility and points of access. This way faculty can distribute the latest article that came out before a class starts. The point of access thing is on the students' side: if it is 2 am in morning, a student isn't near their room and has left their textbook, there isn't much they can do. One thing I want to be careful about is saying it's bad for people to print things. Printing things multiple times or printing textbooks is a bad idea.

Kyle Johnson asks has there been any look at if we can shift large print jobs to somewhere else?

Ed Gomes says there was a possibility that Devil's Duplicates could be an ePrint station, and a job of a certain size could be sent there. We do have the ability to redirect a job of a certain size now to a large volume print queue.

Kevin Davis says the one challenge of this is that articles will be posted online in chapters or sections, so if this to be deemed worthy it will be the responsibility of faculty to keep things in certain places.

Deborah Jakubs says there are restrictions in e-reserve size in place at the library, such as that staff are not allowed to digitize more than a certain part of book.

Robert Wolpert says the bigger picture is any time you distort the economies of a market, people exploit it. We want to do that a little bit, but I think we're seeing what happens when you go a little too far. There may be some point at which we impose some cost for printing. The cost could be inconvenience, like having to walk over to Devil's Duplicates to get something, some sort of market mechanism that is soft so that most printing doesn't trigger it.

Chris says I think it's more about the faculty. If I were a department head I'd want to know if people in my department were assigning students in ePrint, or if faculty were still photocopying on department machines for all their students.

Tracy Futhey says I think we're trying to figure out what is the right balance. I think we should ask Daron to weigh in as well. My guess is that the printing system we have now is extremely convenient for me when I'm first in line to get printout, but extremely inconvenient if I'm behind a person printing out a 300-page book.

Daron Gunn says I've thought about just having one printer for high-volume printing at a print station. But, for example, there are two printers in the Well, and often one of them is broken. If maintenance weren't an issue, that would be ok.

Ed Gomes says there is another issue that students have an attitude that they want to have it right now. Students can wait to print something anytime in the next 12 hours. There are many things the system can do so students don't feel stuck.

Greg McCarthy asks how common is it for universities to charge for printing? Anytime you make something free and sit around asking why is it being used so much, it's like bandwidth.

Ginny Cake says when we first started talking about ePrint years ago we were one of the few Ivy+ schools without a printing quota. Most schools had a quota that students would not reach. After a student would reach a quota, they would have to pay.

Kevin Davis says the other school with no quota policy is MIT, where they have an honor policy of printing only one copy for yourself of something. Ken Hirsch says it always comes down to same thing at the Law School : printing is a cost, and the question is who has the burden of the cost and how can that burden be shared. The crux is ultimately whom or how many of whom is going to bear the cost. The tough part is figuring out how to share that burden.

Molly Tamarkin says Nicholas School students come to me about twice a year very interested in a printing quota; they are not liking the waste of paper they are seeing. I think you'd see some support for a quota.

Mike Pickett says we do have breakdown, for example, of some faculty who are also very "active users." We could set a quota if we want to and have significant impact on heavy users.

Robert Wolpert asks can we have someone look at what quotas are at the other Ivy+ schools?

Mike Pickett says yes.

Daron Gunn says we can't forget in talking about sending large print jobs to Devil's Duplicates, students are doing work later and later after normal business hours, we still need to cater to "I need to print 300 pages at 1 a.m."

Tracy says I think there are many similarities with where we are now with this where we were several years ago with bandwidth in dormitories. The approach we ended up with in bandwidth was good: students came to us and said "we don't have fair access to bandwidth, we'd like you to implement a system so we can have fair access," and we almost ended up with almost an opposite of a bandwidth threshold. I don't know if there is an analogous approach that we could take, but my preference is something proposed by students or faculty rather than the administration saying here's what you can/can't do.

Daron says what also is good about the bandwidth arrangement is that it left the opportunity for people who need that bandwidth to get it.

John Board says so the action items are to do a survey of peer institutions and get data on our own users.

Tracy said she would also like to see someone get with students to empower the student body to make rational decisions as to what would they like to see.

III. Franklin Center IT Voucher Pilot - Mark Olson

Mark Olson says I've brought two of my colleagues from the Franklin Center , Jason Doty and Brett Walters. I'm going to talk about the Franklin Center voucher experiment. For those of you who don't know, the Franklin building is about 22 units with a central support structure. It's small enough to be a good place to try out experimental ways of doing things but also active enough to be rigorous. We experienced intense growth from 12 to 22 units, there are now between 16 and 18 classes per week taught, and the number of events has grown from 250 to 700 a year. Like everywhere else, technology is permeating more and more of the daily fabric of activities. Out of this context, we were doing some audits of our operations, and we were given a suggestion by Dean McClendon to try out some kind of "Monopoly money" goals to place symbolic value on IT services at the Franklin Center . It was also very important that we stress this is a symbolic economy. We wanted to have a control mechanism, because we grew from a small group of folks and as we have gotten bigger we've gotten more beauracratizing and made folks realize we have finite resources. We track desktop support through the help system, so we had a vague sense of where resources were going, but we wanted to get a better sense of that. The last goal is the "is anyone else in same boat" approach, is this scalable?

We created an internal economy, called vouchers. We based it on hours we are available in a workday and started subtracting out commitments to committees. We bracketed out desktop and classroom support since we didn't want units making decisions about supporting these central operation things. We presumed 100% of excess time could go to everyone else. We distributed vouchers and then distributed a fee schedule based on our rough estimate of how much things cost. We revised the schedule as we progressed through the semester based on our learning. Costs were

indicated in “ Franklin Center dollars.” We turned them loose to the end of the fiscal year and let folks spend it. It was a pretty interesting process.

About 37% of vouchers spent were on Franklin Center units themselves. Event support dropped precipitously (we had lots of training, so people did their own tech set-up). The Franklin Center itself is often called upon to execute special projects either across the 22 units in building or for Arts & Sciences, so a very large chunk went to that. In reflecting on it, we do note that the system had the desired effects of controlling demand. None of the units went over. It cut down on doing something because it is there; it got people to think if doing something would be appropriate.

We do know that there is a leak in the system: the Franklin Center is used more and more by the rest of the university. We don't want to curb that, but we do need to be able to account for it. How do we convey that although it's a free service, it does have value? I know Arts & Sciences has model in which they use mock real fees but don't actually charge you. It's not as much a voucher question as a procedural question. There is some concern the system had a negative impact on innovation. As far as the Franklin Center trying to push innovation, last year we didn't do as much as in past years. That's something we want to address for next year, trying to fix some things, having special allocations for special projects, and definitely allocating more of our time to keeping track of all of this. In the process we have an advisory board at the Franklin Center for IT issues and we will decide how to set a certain amount of vouchers for innovation, link them to the Duke Digital Initiative for non-class use, and how podcasting works with a lecture series. We are about to re-deploy another experiment. I'm not sure how it scales and whether it's translatable.

Mark says we didn't have a formal assessment, though we do have a very vocal group and no one said they had big problems. I think most people are sensitive to the fact that we're getting bigger and are sympathetic to the fact that we have limited resources.

Shalesh Chandrasekharan asks do you have sense of budget and if you met that?

Mark says we came up with an allocation by starting with the hours we work in a day, keeping in mind there was some interest in pulling this back to real time.

Robert Wolpert says long ago Pratt had a voucher system for allocating mainframe resources. Initially people took it seriously and it worked, but over a short period of time people became cynical about it. What was insidious about that particular program was that the mainframe was a 20-year investment when technology was moving much faster than that. We didn't buy mini-computers because the mainframe was free. A voucher system might drive people to make irrational decisions.

Kyle Johnson asks how did you divide the vouchers up?

Mark says we allocated them on an equivalent basis, and we made it explicitly clear that if they're not going to use them, they should share them. We're happy people didn't spend them all, because we didn't give ourselves enough wiggle room.

Mike Pickett says about 15 years ago they set up a similar voucher system in the medical center. It was seen as benign, but too complicated to navigate. Did you look at other places that used this to see if it's sustainable?

Lynne O'Brien says I'm thinking back to the last strategic plan which wanted to encourage the broader use of technology. At the same time we don't want people to be wasteful. If we say we want people to stop and not let people innovate, which especially at first may be wasteful, I'm not sure that it fits with our strategic planning.

Mark says we're aware of that. I feel particularly in my history with the Franklin Center as we've grown and had to beauracratize, I'm very concerned about that. But we can't have complete staff burnout because of 700 events.

Molly Tamarkin says one thing that is useful with this kind of system is you can actually see where you are growing the most, to know where need resources the most or quell perceptions that one group is taking up all your time. It could be used strategically to determine how to develop your department's plan.

John Board asks did you notice any change in requests on the classroom side that indicated culture change?

Mark says the only culture change we noticed was more self-sufficiency on the part of the faculty. We shifted to a role we want in terms of troubleshooting. We had lots of questions about desktops and how to manage that. I think having numbers as an institution as to what appropriate desktop numbers are would be helpful.

Mike Pickett says within the university there is a lot of variety. I saw a recommended desktop number and it was 51.

Tracy Futhey says this may be a topic to return to over the course of the fall. As we proceed with the strategic planning piece, how do we provision for support? If there is a good model out there, let's get that out there and know what it is.

IV. Updates on Library card catalog - Ed Gomes

Ed Gomes says the online catalog had a hardware failure that manifested itself Sunday morning at 10:45. It was first diagnosed as a hardware failure, then a failover failure, then both, so between Sunday and Monday we spent a lot of time talking with folks from Sun and Ex Libris, who manufactures the software system. We were able to get some parts shipped in from Charlotte , and once we replaced the parts on the hardware, we had to properly get the software to talk to the hardware. The discs were showing up as properly configured, but a volume manager wouldn't recognize them. We loaded a patch recommended by Sun Monday night, and at Tuesday morning at 8:00 we got everything put back together, loaded up the system, and the database was corrupt. Even though our procedures were in place to do proper backups, there was some oversight in terms of managing the process and the eventuality is that the most recent copy we had of the database was from July 21, 2005, so everything from then until now is gone.

We spent a day and a half working with Ex Libris and folks from OIT who graciously offered their services rebuilding the database and getting the system backup online. In the interim we were running a failover system so people could search, even though we had removed the part where someone can check out items so they couldn't check out books that were not there.

After the back and forth of bringing the system online we went back into production today at 2:30. All services are restored. We are currently loading all of the patrons onto the database system. That data load should be completed by about 9:00 pm this evening. We have loaded course enrollment into the system, and there will be some hardcopy reserve items that will have to be manually added back into system. Circulation transactions are gone, so one thing we are dealing with is people logging in and seeing items out that have turned in. We are suggesting that patrons physically bring items in so we can get them back into the system.

We briefed the staff yesterday in several meetings. We started to try to identify sources of data in various locations on Monday afternoon knowing this could possibly be an issue. We are continuing to search those sources of data to retrieve as much information as we can. We have some paper records of acquisitions and invoices processed during that time. We've got a lot of work ahead of us.

John Board asks is there any hope of classic data recovery?

Ed Gomes says we've had the OIT Oracle experts look at it, and its not possible. The real kick in the pants is when the hardware failed, our latest backup of the database was one hour from completion. A lot of unfortunate things happened, but the library folks have been very good about it. Some faculty have been complaining about not being able to access E-Reserves, but they are now aware they can access them through Blackboard.

Robert Wolpert says it looks like we have had two high profile disasters lately. We realized in our department, on a small scale, we were no longer making backups of what we should be. It seems like a pattern.

Ed Gomes says it's not that people become complacent, but you get a backup strategy and check it repeatedly and eventually say, its working so I don't need to check it as much.

George Oberlander says there really is no good way to check backup. A good way is to have another machine.

Ed Gomes says we had a backup box. We were following a model where we were writing the backup to disc and storing it for seven days. We would also backup that backup folder to tape. What happened is that the backup directory got dropped from the tape rotation. That system was running properly until July 21 st , and then it stopped functioning properly and it was not picked up by the monitoring system. It wasn't that procedures weren't in place but that things happened and were not followed up on in terms of monitoring. You can't assume that because something works correctly for a year going it is to keep working correctly. Now OIT is managing our Oracle support. There was some delay in bringing the system up because we wanted to bring up their mechanism for backup, system failure, etc.

George Oberlander says another thing to do is to backup with a different product, to use another vendor's product to do an image.

V. Email/Webmail update - Michael Gettes

Michael Gettes says the Webmail environment OIT provides was running IMP Version 3 and we switched over the weekend to Version 4. Over the course of the last year we were experiencing various anecdotal performance problems. Especially last spring things came to a head, and fortunately Version 4 was coming out and they also redesigned a particular area that would solve some performance issues.

One new feature is shared folder support; rudimentary filtering capability (there are some issues but it is there); one of nicest features is when you log in, on the top bar it shows you your quota usage in percentage and actual numbers. We provide Webmail services to ten other schools within Duke through the same interface.

One of the neat features is connection caching. In Version 3, every time you rendered a web page it would involve a log into the IMAP server, then disconnect. In Version 4, you log in and it keeps the connection open for some period of time, and it has in impact. We ran into a performance problem with all of this in production that didn't appear in the testing, especially with CPU utilization. We found out that MIT had the very same problem. To make a long story short, we decided to throw more hardware at it. We were fortunate that we had new cluster nodes and we added those. We are now buying seven addition nodes and trying to get that done as quickly as possible. While we can't prove it with the numbers due to the nature of the way things were, we think that about 75% of mail access is through Webmail.

Robert Wolpert asks is it easy to understand why CPU usage jumped so much?

Michael says we couldn't figure it out and neither could MIT. Now what we're seeing is that it went from 7 log-ins per second with Webmail to an average of 1 per second. That had a significant impact not only on the Webmail environment but on the Web in general. We said that by this time we would be on new disc arrays supporting larger quotas, but we ran into problems with configuration - the vendor got it wrong, not us. We're hoping for some time in October to bring the new disc on-line. To get to the new discs will require some down time, and that will also allow us to put Webmail on the new machines.

End time : 5:30 p.m.

Duke Office of Information Technology - www.oit.duke.edu - (919) 684-2200 - help@oit.duke.edu