UNIVERSITY OF CALIFORNIA ACADEMIC SENATE
UNIVERSITY COMMITTEE ON COMPUTING AND COMMUNICATIONS
Minutes of Teleconference

April 20, 2011

Attending: Joel Primack, Chair (UCSC), Anthony Joseph (UCB), Ida Sim (UCSF), Jun Li (UCR), Sreenivas Jammalamadaka (UCSB), Brett Stalbaum (UCSD), Felix Wu (UCD), Maurizio Forte (UCM), Lawrence Lin (Graduate Student Representative, UCSF), Robert Anderson (Academic Senate Vice Chair), Stephen Lau (Information Resources and Communication Director of Policy), David Ernst (Associate Vice President and Chief Information Officer, Information Resources and Communication), Warren Mori (Co-Chair, Shared Research Computing Pilot), Brenda Abrams (Policy Analyst)

I. Welcome and Announcements

Chair Primack welcomed the members to the meeting.

II. Consent Calendar

Action: The minutes were approved.

III. Consultation with the Office of the President

- *Stephen Lau Systemwide IT Policy Director, IR&C*

Director Lau joined UCCC to discuss the UC Privacy and Security Initiative. Two aspects of privacy are the issues of civil liberties and data protection. This is an 18 month initiative launched by the president to look at how UC views privacy and information security, to consider best practices and determine whether UC's policies need to be updated. Current policies and rules were written prior to things like videoconferencing and shared resources so there is a lot of conflict with new laws. Faculty, staff and students have an expectation of privacy unless there is a justifiable reason or consent has been given. Emails and web traffic are not monitored. New regulations ~~that~~ make UC responsible to ensure that medical information is not leaked. In industry, Social Security Numbers are monitored to prevent their being released accidentally. UC does not monitor these things so there is a risk that this information could be leaked, but monitoring runs counter to UC's culture. Two issues include the electronic aspect in terms of monitoring for privacy and the civil liberties aspect where people feel they are being watched. UCB Professor Chris Hoofnagle, a privacy expert, spoke to the steering committee members. Most academic institutions are just beginning to explore this problem. A member of the UCLA privacy board also spoke to the committee. This board looks at these issues on a campus level and UCOP is exploring whether there could be a similar board across the system and established at each campus. One situation that occurred was that a computer science researcher wanted to conduct search engine research and monitor network traffic to see how people were utilizing search engines such as Google or Yahoo with the goal of optimizing them. UC's policies prohibited the researcher from conducting this type of monitoring. The researcher contacted the

search engine companies himself and was able to purchase the data, thereby circumventing UC policy.

Discussion: One question is whether Google sells information with or without identifying personal information. The UC medical centers have a policy that all laptops and desktop personal computers must be encrypted. The problem would be that someone could decrypt this information and distribute it. If the goal is to keep information private, this will not be achieved through monitoring email. It may be very difficult to develop one policy that addresses privacy and security in all situations. More training might be a better approach so individuals understand how information needs to be handled in order to prevent accidental release. The steering committee is discussing the how much should be monitored and the risk UC is willing to take. Monitoring email may not be the best strategy to prevent the malicious release of information. Prevention of unauthorized access needs to be in place. California's public information act allows anyone to request information from UC.

In a recent case, someone requested a University of Wisconsin professor's emails under the state's open access law. The University's Office of General Counsel determined that anything protected by law would not be released and the clause that it was not in the public's best interest to release the professor's emails related to his research. UC has not received a request related to faculty research or for political reasons although Vice Chair Anderson indicated that in the last month at UCB faculty have discussed requests for information from faculty. Director Lau stated that if any UC faculty ever receive a request under the public information act, it should be brought to the attention of campus general counsel. It is not clear at this time whether UCOP would follow the precedent set by the University of Wisconsin. A member asked if it would be better to move campus email to an external provider but it is not known if someone could make a request to the provider directly. One reason to use Gmail would be that a court order would be required to get access which would necessitate more effort. In this scenario, someone could ask UCOP for the information and UCOP would ask the faculty person to provide it regardless of where it is stored. Every campus and department has different policies with respect to when emails are permanently deleted from the servers. UC will need a strategy that is directly responsive to the criteria in California's public information act. There should be a discussion about what is and what is not covered by the act. Director Lau reported that OGC is very sensitive to the impact of any decisions or policies on faculty. UC should ensure that transparency is not jeopardized so any policy needs to balance the need for privacy. UC's policy needs to ensure that academic freedom is protected and the public information act may need to be changed to address this.

IV. Campus Gmail

According to Chair Primack, the discussions with Google about moving to Gmail for faculty and students have hit a stumbling block. [However, this was subsequently resolved.]

Discussion: Students at UCD are using Gmail, and the students are reportedly very happy with Gmail. There ~~re~~ is a small pilot at UCD with a few hundred faculty and staff. ~~The students are reportedly very happy with Gmail.~~ The goal of immediately transferring all faculty and staff has been postponed. At UCD, there are 100 different email systems being used by different

departments. There was a push from students and staff to use tools like Google documents, calendaring, and other applications. The campus is concerned about Google mining information if it is used for everyone. Other concerns are that emails could be leaked out and that the servers are outside of the US. Microsoft Office 365 guarantees that emails will be hosted in the US, whereas Google would charge UCD to host the emails in this country. Faculty doing certain types of research will be restricted by ITAR from using email providers with servers outside the US. UCD will continuously monitor whatever system is used if the email is ultimately outsourced. A member commented that it is important to have flexibility to accommodate different types of usage. UCD's campus--wide email will exclude the medical center but will share feedback on the experience with the center's leadership. Office 365 allows users to have emails hosted locally and the basic service is free. This system also offers larger mailboxes than Google and it also allows for larger attachment sizes. Chair Primack reported that the Chinese government claimed that there was a break-in to Gmail. The advantage to Google for hosting faculty and staff email is not obvious. One question is whether UC or Google would be liable for the unauthorized release of emails.

UCSC switched all students to Gmail last year and it was a success. UCSD provides and manages a centralized exchange server to students but this has not resulted in notable cost savings. UCSB has not changed its email system. UCB runs its own in-house email and it is an inexpensive solution. UCR students have switched to Gmail. UCM has not switched to Gmail and there is no plan to do this in the near future.

V. Shared Research Computing Pilot

- *Warren Mori, Co-Chair, ShaRCS*

Chari Mori joined UCCC to provide an update on the pilot. It is important for UC campuses to work together to share resources and share expertise across the system. This project will allow UC to pursue large multiple PI grants. Local data centers were not utilizing resources effectively.

Discussion: Chair Primack described several issues with the pilot. One is that many of the PIs involved in the pilot already have access to bigger systems. Chair Mori reported that turnaround time is better for his research. The pilot will be evolved into a shared cluster. New users will have to use grant money to buy nodes that will be added to the center. Chair Primack commented that NSF may approve of researchers using the same cluster. Fourteen PIs responded to a survey and IR&C will provide UCCC with all of the responses. Users reported that the cluster exceeded their expectations, that the project should continue into the foreseeable future and that they would participate when it transitions into the new shared cluster model. The northern cluster is used about 80% of the time.

Procedures are in place to prevent any one group from using all of the resources and to limit the length of time for a run. There is large variation in the number of nodes requested by run. The initial set of PIs in the pilot was chosen because they had experience with large machines so the clusters could be effectively utilized. Additional resources in order to provide expertise and consulting would be required to help individuals who lack experience with this type of machine and this is being discussed now.

Roughly the total cost over three years for a current node is $12,700. The node is $5000, electricity is $1500 and there are also system administration costs. The PI pays upfront for the node. The expected lifetime of each node is three years and if it lasts longer, the price will go down. UCOP, the campus and researchers contribute to the electricity. A comparison between the cluster and Amazon's system has been made and the UC cluster costs $.04 per cpu hour. This will need to be reassessed. With Amazon, it is cheap to input the data there and expensive to export data. Chair Mori does not think that the clouds can compete with UC in terms of cost. Other leading institutions are providing the resources and facilities for faculty of the same magnitude and are moving toward shared clusters. The data centers are state of the art facilities. The UCB representative reported that Amazon's cost per cpu is $.10 before a discount for a bulk purchase. The depreciation cycle cost has probably not been taken into consideration when UC is determining the cost. Chair Mori indicated that researchers who participate in a shared system get a lot in return. UCOP should explore what could be purchased if the cloud is used, and a member noted that Amazon is willing to build different types of nodes as requested. If UC has to absorb another significant budget cut that results in stopping this pilot that will be very problematic for the participating PIs. UCCC would like more direct comparisons with commercial providers which will be especially important when faculty have to start paying for the cluster. It is not clear if certain funding agencies will grant funds to use the cloud whereas researchers are more likely to provide funds to use the shared cluster.

VI. Consultation with the Office of the President

- *David Ernst, Associate Vice President & CIO, Information Communications and Resources*

AVP Ernst reported that UC is on the verge of signing a master contract with Google that would enable campuses to take advantage of Gmail for students, faculty and staff at no cost to the campuses or users. Two issues have been outstanding in the negotiations. These are ensuring that Gmail is accessible to people with disabilities and Google being willing to guarantee that data stored for their users would be within the boundaries of US or jurisdiction of the US court system. A conference call with Google leadership will occur in the next week. UC is prepared to sign the agreement if the guarantee is not included but there will be serious caveats about the risks.

AVP Ernst reported that the new LBNL building would open in late 2014 if it is built at LBNL but if it is built on another site it would open in 2015. It is not clear if the facility will be in a new building or in an existing space. UC needs to determine its longer term plan to provide research and general computing resources in secure regional facilities at a lower cost than would campuses would have to pay. Regional UC facilities may be a transitional step toward doing most of the computing in the cloud instead of at UC. The supercomputer centers will probably be full in two years. The center at UCSD is not a long term solution for UC. Commercial co-location is being discussed and CENIC may be the middle man between Amazon and the CENIC partners. There is interest at UCOP to move toward regional computing outside of or at UC sites. By the end of the calendar year there will be a proposed strategic plan that will be evaluated by the system. There is willingness at UCOP to invest in strategies that will offer savings in the longer term.

There is a plan at UCB to begin charging for electricity for computing clusters on the campus. AVP Ernst indicated that a plan is need to enable buildings to be metered so individual departments and faculty can be held responsible for their power consumption. If departments are able to keep the funds saved, faculty behavior will change in order to generate even more savings. There is no plan to implement strategies to reduce energy usage and costs in the next academic year.

Chair Primack stated that the ReadyTalk videoconference is an improvement over iLinc but both are far less sophisticated than the Evo system. There will be discussions by ITLC beginning in May about updating the inventory of the systems being used at UC. Identifying the basic features of a system that people would want to use will be the first step. AVP Ernst indicated that ReadyTalk is one of the most widely used systems at UC.

There is currently no strategic plan for digital data storage although there have been discussions about this. At UCM, faculty are asking what the plans will be for backup storage and the IT department is asking faculty to pay a fee which is not possible with certain types of grants. It is not clear who is responsible for managing the data. The campus charges $5,000 for ten terabytes. It is not clear what UC will do if there is another significant budget cut.

Discussion: ITAR covers control of export information and if the information is taken outside of the US, this is an export event. Due to ITAR it will be important to get the guarantee from Google. There are companies that conduct energy audits and pay for the capital costs of replacing inefficient infrastructure such as lights. This is in exchange for the institution paying this company the amount of its existing electric bill over several years so the company recoups its costs. According to AVP Ernst, a task force will be established to explore how these companies could be used at UC. Under the new scheme at UCB, certain departments will be penalized because they have reduced power usage over the past several years in certain buildings. NSF proposals now require a data management plan. The California Digital Library is offering a free service to help faculty develop these plans. The CDL may help faculty connect with data storage facilities. UCOLASC has discussed the CDL project and UCCC should discuss this matter with that committee. UCCC could suggest features to CDL that would be valuable. NASA is also now requiring data management plan and developed protocol years ago for the management of data. Foreign countries are ahead of the US with respect to data storage.

VII. Budget Cuts

It is not clear what will happen with respect to IT in order to manage the budget cuts.

Discussion: Currently faculty have equipment in closets so UCB is looking at creating a number of tiers of data centers in terms of reliability, availability and power utilization efficiency. There would be incentives to move equipment to the campus data centers. The campus data center has not been very reliable and it is close to the limit in terms of thermal cooling so new facilities would need to be built. UCSC is also running out of space at its center and there is a pilot project to remotely place data at UCSD. Shipping containers are being used as pods at UCB. A significant investment by UCSC would be required for their remote facility to be used but these investments would result in savings in the future. Vice Chair Anderson noted that there are loans

through UCOP that might be used by campuses for this type of work. According to Chair Primack, as much money will be spent on electricity, including cooling, as is spent on hardware over the course of three years.

VIII. Member Items: Major Campus or Systemwide Issues

UCSD: The representative reported on the online education pilot project. He participated on the selection committee. Overall, he feels the pilot is being implemented well and that selection process was done fairly. There was an effort to include all of the campuses and a variety of disciplines. Online education will not save money but provides an opportunity to generate revenue from students outside of California. There is an effort to ensure quality. UCCC should monitor the current status of the project. Some courses will be ready to launch next fall. A vendor will be selected to develop the online infrastructure. A RFP for a consulting company will be released next week. UCCC should identify any associations between the companies and anyone at UC, e.g. the Regents.

Educause provided a grant to fund some of the courses ready to be offered soon. There is a C3 loan available for cross campus collaborations. A total of $6.9 million in loan funding is available and there will be controls in terms of how this money can be borrowed and spent. A concern is how the money will be repaid if the project fails. The representative will forward materials to the committee for review. It is not yet clear how the courses will be reviewed and approved. Online courses could make more specialized courses more widely available although they might also eliminate introductory courses thereby decreasing the number of faculty in a department.

Meeting adjourned at: 3:15
Minutes prepared by: Brenda Abrams
Attest: Joel Primack