

Designing and Enacting Change: Penn Restructures Computing Services

CAUSE '96
December 5, 1996
San Francisco, California

Mario Moussa, Ph.D.
Associate
Center for Applied Research
Philadelphia, PA 19104
mmoussa@mail.cfar.com

Linda May, Ph.D.
Director, Planning
Information Systems and Computing
University of Pennsylvania
Philadelphia, PA 19104
may@isc.upenn.edu

Lynn Oppenheim, Ph.D.
Vice President
Center for Applied Research
Philadelphia, PA 19104
loppenheim@mail.cfar.com

Noam Arzt, Ph.D.
Executive Director, Administration
Director, Information Technology Architecture
Information Systems and Computing
University of Pennsylvania
Philadelphia, PA 19104
arzt@isc.upenn.edu
URL: <http://nextb.dccs.upenn.edu/noam.html>

Abstract. Penn is restructuring the way it organizes, staffs, and funds computing services. Seven pilots are testing a new model developed by a campus task force and vetted throughout the institution. But a restructuring is just a diagram on a piece of paper until behaviors themselves change. This presentation shows step-by-step how the structured, participatory process of designing the new model was itself a way of creating—and trying out—desired new relationships and outcomes.


Designing and Enacting Change: Penn Restructures Computing Services

In this paper we tell two intertwined stories about Penn's recent restructuring of its computing support services. Both stories have a point to make. The first is about the restructuring process, and offers an alternative to what might be called the "classic four-part reengineering saga": analyze an organization's work processes, redesign them, put them back together again, then get people used to the changes. In our case, by contrast, the very start of the restructuring process ushered in new ways of working. The second story is about the outcomes of the restructuring. As you will see, Penn has placed the user at the center of its computing support model, which is distinguished by flexibility, trust across academic and administrative divisions, and a willingness to learn. We want to show how these active qualities are essential to the ongoing success of the user-centric model and were anticipated by the design of the process itself.

We contend that an organizational model—restructured or otherwise—is nothing more than a picture on page unless it is enacted, daily, by people in their professional roles. In this respect, behavioral change is at the root of organizational change—and effective restructuring efforts. The carefully designed participatory process we describe in the following sections was a means of creating the behaviors required by a computing support structure that would satisfy Penn's diverse needs. We begin with i) an overview of the organizational model Penn created in response to those needs, then move on to ii) an account of the process by which we arrived at it, and conclude with iii) an examination of the workplace practices that continue to affirm and reaffirm the model's fundamental structure.

Penn's New Model

The computer user is at the center of Penn's new model. Each person should have a local computing "home" and take all computing questions there. Beyond this circle of frontline support are expanding circles of secondary services, provided by the schools and by the central computing group, Information Systems and Computing. But the map of services is irrelevant to recipients: the frontline support person navigates that landscape for them.



The model clarifies a division of labor that lets the schools and the central computing group each devote more effort to the things it does best. All frontline support is at the local level as close as possible to the user. The central computing group, in turn, concentrates on networking, core business systems, data administration, standards, and services for local providers of computing. The model requires trust that other people will do their job—and is predicated on the belief that doing a few things well is better than doing a lot of things poorly.

Two new strategies make the central computing group more accountable. The first is to run the network as a regulated public utility. The second is to sell services where markets exist. These small businesses are likely to remain a small part of Penn's overall economy, but they're an important Trojan horse, exposing the institution to the discipline of the market and the pressure to please the customer.

A third force (not the center, not the individual schools) focuses energy at the University level. A few cross-University processes will be funded directly and

managed across traditional organizational boundaries. These process teams build flexibility into the system. Subject to sunset laws, they have to prove their value in changing circumstances.

Appendix I describes the model in more detail, as does the project's Web site (<http://www.upenn.edu/restruct/>).

Penn's new model is well into the pilot stage. Appendix II describes seven pilots that are testing aspects of the model.

The Restructuring Process

We designed the new model in four very public, very structured retreats over the course of a semester. We wanted to send a message that endless meetings were not in store. We know there are other ways to design a new computing support system, but we want to highlight ten features of our project that might be worth borrowing.

Time-bounded urgency. Our sponsors gave us just one semester to come up with a new model. We knew that their patience—and our contract with our consultants, the Center for Applied Research—would run out at that point.

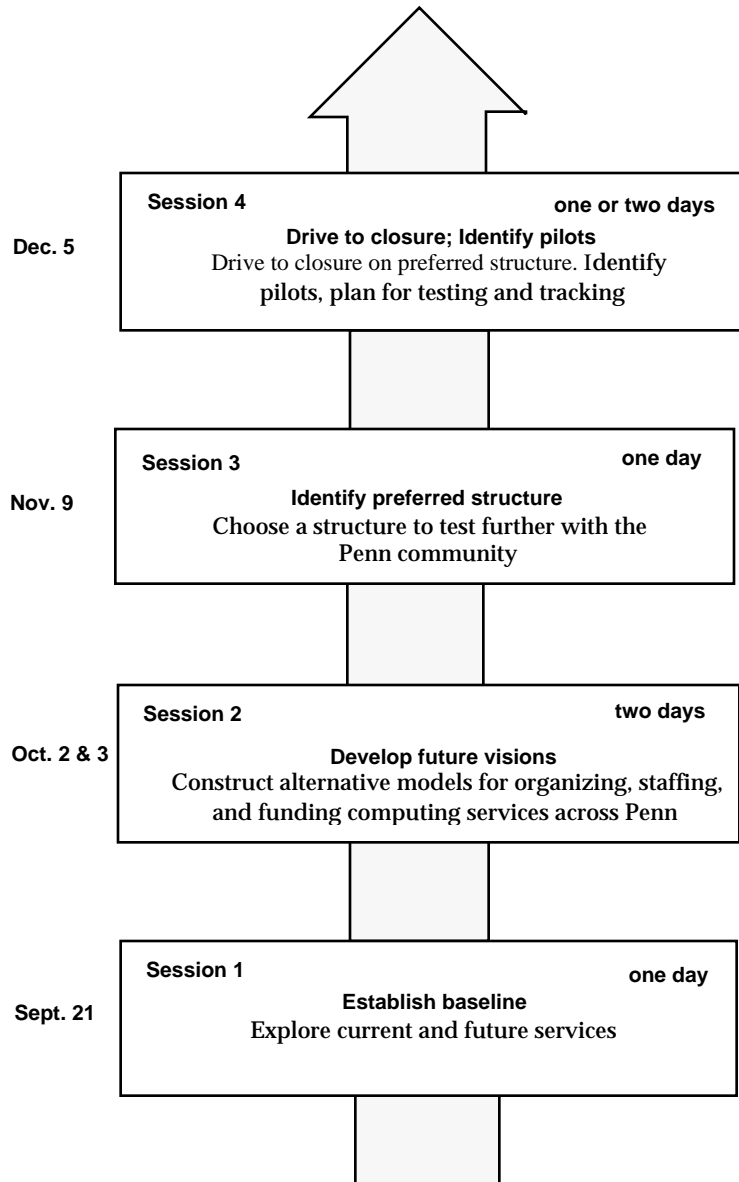
System in the room. We recruited twenty-eight people to the task force. They included users of computing as well as providers, people from the schools as well as the center, and, strikingly, eight faculty members along with administrators. The task force was co-chaired, in fact, by Penn's CIO and by a full professor of Classics.

Structured retreats. But "system in the room" is easier said than done. A twenty-eight member task force is usually too big to get any work done. So we structured the retreats as a combination of small-group problem solving and full-group evaluation and synthesis of the work of the small groups. Assignments were handed to three subgroups—Benchmarking, Models, and Communication—who worked hard between retreats.

Real decisions, real time. We used the public nature of the retreats. The task force found it energizing to be making real decisions in real time. And we shamelessly exploited people's fear of looking bad at the next retreat if they didn't do their homework. If circumstances don't give you an obvious set of deadlines, we recommend, in fact, that you create public events to build urgency and tension into the system.

Project-on-a-page. The discipline and communication value of distilling the whole process into a single image is substantial. We used the image we called our "project-on-a-page" over and over (below and see Appendix III for our more complete handout). At the beginning of every retreat, we'd say, "You

are here." When people wondered if the process was going to come together at the end, we'd come back with, "You are here."



Honoring the doubts. We respected the University culture of criticism and rigor but didn't let it paralyze us. In the very first session of our very first retreat, for example, we aired the doubts of the participants. We called the exercise "My Skeptical Friend" and asked people to interview each other:

Imagine a valued and skeptical friend asks you a tough question about the restructuring process. What would that question be? And what can this group do to keep your friend's worst fears from materializing?

We heard questions like these: "Is this really about firing people?" "Won't it be all talk and no action?" The phrase "my skeptical friend" entered the lexicon of the group. Throughout the process, people would surface their own doubts by mentioning their skeptical friend.

Forced, but provisional, decisions. As the story, below, of the winning model will show, each of the four retreats required decisions at specific points—but decisions that could be changed and improved. In this way, the group began to feel comfortable about moving forward.

Divergence then convergence. We designed a process to encourage divergent thinking about alternative models. Since comparison yields insight, we didn't want people to come to closure too quickly. At the second retreat, we charged four small groups with coming up with four different models for computing support. They could *not* base their models on dimensions Penn usually uses to think about itself: centralized/decentralized; large school/small school; allocated costs/direct charges.

The small groups produced rival models. One was the Process Model where cross-functional teams work across organizational boundaries. One was the Revolutionary Council where citizen patriots from across campus meet to make decisions about allocating resources. A third was the Village Model where all computing support is local. Resisting all efforts at early synthesis, we handed the models to a workgroup and asked it to flesh out each one.

At the third retreat, we tested each of these models against a series of tough cases like exploding demand and need for accountability. At the end of that day, one of the models clearly emerged as the most robust. That model ultimately became the one you see in Appendix I.

Can't learn from what never ends. We built in end points all along the way to allow for reflection and assessment. We debriefed at the end of each retreat, for example. (How much progress have we made? What could we have done better?) The project had a clear end date, which kept us from tinkering endlessly with the design. And this paper itself is an end point, a chance for us, and our readers, to clarify the lessons we learned from the restructuring project.

Pilots as a learning—and clarifying—principle. We were emphatic from the start that we weren't trying to design a perfect model. We knew we would subject the model to a pilot phase, testing it against Penn's gritty reality under the watchful eye of an implementation steering group. The pilots, moreover, (listed in Appendix II) gave the model a vividness that was lacking in the design stage. ("Oh, *now* I get it," more than one person told us.)

The Outcomes

A. New Ways of Working

Having read about the process, you might reasonably ask, "What's now substantively different about the computing world at Penn?" The pilots in progress (Appendix II) illustrate three practical ways the model has changed the way Penn does, and thinks about doing, business.

Bang for the buck. When you are restructuring you have to be able to show that you've made a difference. Penn's new model helps locate bang for the buck. It suggests which pilots will make the most difference—and which ones have to be done *first* or people won't take the model seriously. The Support-in-Residence Pilot, for example, aims at moving frontline support for undergraduates out of a central help desk into the student dormitories. This is a highly visible—and literal—demonstration of a main principle of the model: Every user should have a local computing "home" and take all computing questions there.

The pilot team knew financing would be a challenge, so they looked closely at how money and people were already being used and let the model be their guide to repurposing them. In this case, work-study students paid primarily to guard laser printers in residential computing labs were organized instead into a nimble team that answers students' computing questions and even makes house calls to dorm rooms. With training, the work-study students are able to answer most questions themselves—and know which ones to refer up the line to a computing professional. Piloted successfully in one student residence in 1996-97, the support structure will be expanded to all first-year and college-house residences in 1997-98.

Saying "no." What's really hard, particularly in support, is saying "no." Penn's new computing model—founded on the principle that doing a few things well is better than doing a lot of things poorly—helps people know *when* to say "no" and, likewise, when to say "yes." The most striking new examples of "no" and "yes" are in the area of frontline support. As of July 1997, the central computing group will no longer provide general frontline support except on a contract basis. Schools, on the other hand, are saying "yes" to frontline support. Under the new model, schools are responsible for the frontline support of their own members. They can provide it themselves or buy it from each other or from the center. (The Frontline-Support-for Hire pilot, for example, is developing a business plan for the center's support-on-site program.) Many of Penn's schools already operate under this principle, and others are beginning to fall in line. The schools, moreover, are beginning to distribute their own support more locally. The School of Medicine, for example, has just notified its many academic departments that frontline support will become a departmental responsibility. All in all, support is

moving much closer to Penn's computer users—and at different speeds and in different ways, which makes sense in Penn's decentralized environment.

The big picture. The model broadens the picture by helping people understand relationships across the system. Most of Penn's public computer labs for students, for example, are run by the individual schools. Some labs have up-to-date computers and others don't. Software also varies from lab to lab. As part of the New Learning Spaces pilot, Penn's CIO decided to take part of his operating budget and set up a matching-fund program to help labs replace their equipment on a four-year cycle. That was a novel idea: giving away money to improve someone else's facilities. But Penn's CIO believed the savings in support costs to the university—with a more standard and up-to-date collection of lab machines—would far outweigh the funds he freed up. And the model made this politically acceptable as well as achievable in practice.

B. Ongoing Guidance: The Model in Action

By now we hope you have gathered that the model is a living structure that provides almost daily guidance to key people in the computing support system. How?

Model as theory. First, the model is a theory, not a rule book. To us this means that a few principles, not pages of instruction, guide action. It's an example of what CAUSE '96 keynote speaker Meg Wheatley would call "unity in essentials." The model sets out essential principles and allows freedom in the nonessentials.

At least fifty decision makers at Penn make significant choices about computing resources. In such a decentralized environment, the model provides coherence for these choices. When the Executive Vice President wants to put in place a financial accounting process that has computing implications, the model tells him what he needs to do to be a good citizen. When a Dean decides to hire additional staff, the model says where, as a good citizen, she should put those staff and what she should ask them to do.

Guide to decisions not yet made. The model guides decisions that have never been explicitly considered. Take again the example of Penn's public computing labs. The model says that standards are important—but standards, as we all know, are very hard to enforce. Penn's CIO's decision to give away matching funds as an incentive for standards was an innovation, a creative application of the model.

Growing body of "case law." The model, then, is continually elaborated in local decisions. It has spawned a growing body of knowledge—a little like case

law—on how to solve problems at Penn. We've put in place a steering committee to guide this early stage of organizational learning. It's at the steering committee that knowledge from independently running pilots comes together and enables us to ask: Do the pieces still fit together? Are the decisions we're making with respect to specific problems still coherent with the model? The restructuring project also has a Web site (<http://www.upenn.edu/restruct/>) that documents a growing body of ideas and solutions. The site follows pilots that are testing different ways of handling frontline support at Penn, for example.

Guide to further exploration. Like any good theory, the model closes off some avenues of exploration and opens up others. Penn's Deans, for example, know it is their responsibility to provide frontline support within their own schools. They don't have a choice about *whether* to do this—but have many choices about *how* to do it. The School of Engineering, for example, has long sustained a flourishing computer support structure that makes effective use of work-study students. The School of Social Work, in contrast, has recently decided to buy on-site support from Penn's central computing organization. The model is flexible enough to accommodate any number of possibilities, even outsourcing beyond Penn. The model, in short, doesn't dictate solutions, but helps people know which problems are worth solving.

Flexibility, with commitment. Any theory must be flexible; the problems that are going to emerge next year are not the ones that even the brightest people can anticipate. As theory, not rule book, Penn's new model allows for change over time and across the University. But the users of the model are also the designers of the model—they know how hard it was to create and they believe the model is about as good as it is likely to get in the short run. They will not jettison the model and start over every time a problem arises.

In sum, Penn's new model provides few specifications—it is not a rule book. It embodies instead a theory, a set of principles, that provides coherence across a range of specific, local decisions.

The Driving Force

You might be wondering what drove the participants in the restructuring process ultimately to come together around a few basic principles. At the root of a very complicated situation was a common notion: nearly everyone involved, providers and customers alike, believed they had to improve computing support themselves—or live with solutions created elsewhere. Those solutions might be forced by the market, by the Trustees, or by the momentum of technology itself, but something was going to force change if change was not actively ushered in. This belief was present from top to bottom—from the executive sponsors, who knew computing had to be restructured in order to play a role in major reengineering projects underway

in other areas, down to the diverse users who depend on computing to get their research, teaching, and administrative work done.

Of course, the history of efforts to change organizations is full of stories in which everybody believes a system needs repair and still fail to agree on solutions that can be implemented. During our project, we made a practice of reminding ourselves that change has to work in the real world where Penn and its people exist. In the real world, those most profoundly affected by restructuring need basic questions answered: Will I have a job when the dust settles? How will my years of experience and knowledge be used, if at all, during and after restructuring? Will my needs be met by a new way of delivering services? Will the restructured organization be built to last, or is it a quick fix?

As Michael Hammer recently observed in the *Wall Street Journal*, reengineering's fatal flaw is a tendency to forget about the people who must live with a reengineering solution. We believe our restructuring effort succeeded because it found the right *kind* of solution, nuanced enough to respond to the needs of faculty, students, and staff across campus. We allowed doubts to be expressed, we encouraged divergent thinking, we forestalled overanxious agreement, and we stopped when we had arrived at principles that could meaningfully guide local decisions.

Penn's new structure for computing services

October 1, 1996

Why a new model? Computing now touches everyone at Penn. Demand is soaring and Penn's investment is sure to expand. Those who use and those who provide computing recognize that support can be improved. In the fall of 1995, Penn's Provost and Executive Vice President appointed a university task force to make computing services easier and more cost-effective for those who use them. The task force has produced a model that will guide organizational change over the next few years. The model doesn't claim to do everything. It doesn't ignore history. It is a way of doing business that gives members of the community the chance to make Penn better and exposes each of us to the costs of bad decisions and the benefits of good ones.

The user. The computer user is at the center of our model. Each person should have a local computing "home" and take all computing questions there. Beyond this circle of frontline support are expanding circles of secondary services—provided by the schools, the central computing group, or by outside vendors. But the map of services is irrelevant to the recipient: the frontline support person navigates that landscape for them.

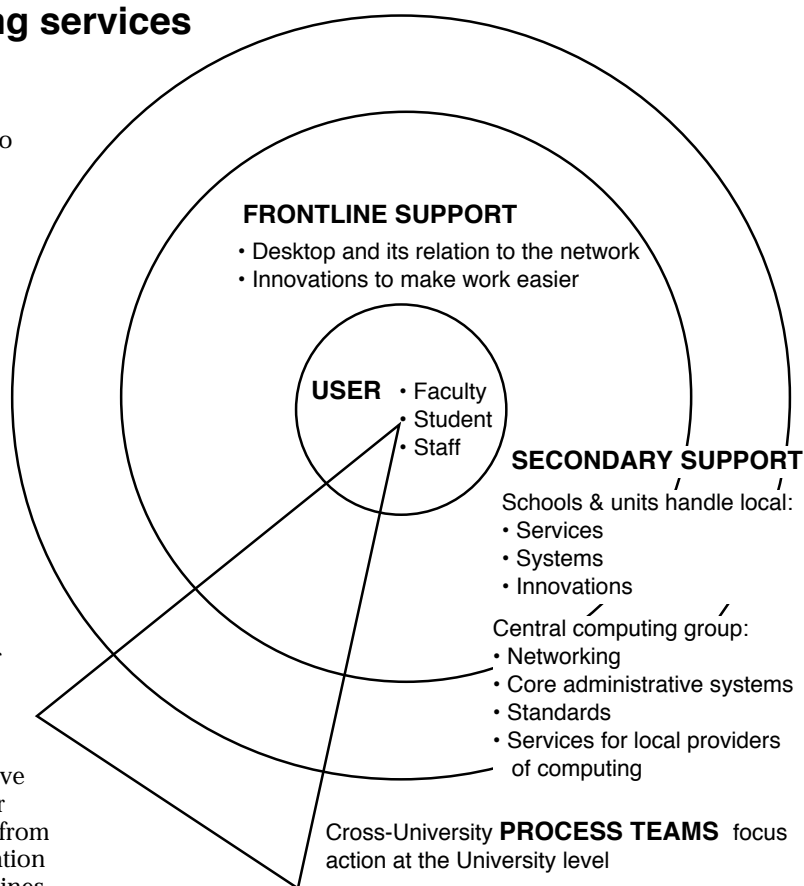
Frontline support. In the model, schools and administrative divisions are responsible for the frontline support of their own members. They can provide it themselves or buy it from each other or from the central computing group, Information Systems and Computing. The task force urges that guidelines for basic levels of frontline support be set and that Penn institutionalize ways of keeping those levels moving up.

Secondary support. Secondary services undergird frontline support and make the whole greater than the sum of the parts. The central computing group will provide a more focused and delimited set of secondary services: core administrative systems, networking, data administration and information security, standards, and services for local providers of computing. Two new strategies help shape secondary services:

■ **Network as a regulated utility.** Penn's network will be run as a regulated public utility—with service-level agreements, campus-wide standards, and a "public utility commission" or campus governing board, to keep it responsive. As a utility, the network will be funded by a mix of allocated and direct charges, with specific funding strategies to be taken up by the governing board. Telephone and video services will be incorporated into this utility structure.

■ **The market test.** The model encourages ISC and others to sell services where markets exist. The intent is to establish, over time, enough of a marketplace to help control costs and encourage a focus on the customer. Small businesses are already found at Penn. Wharton Reprographics is well known; ISC sells support-on-site, training, application development, and integration services. Such businesses are likely to remain a fairly small part of Penn's overall economy. They are an important Trojan horse, however, exposing the institution to the discipline of the market and the pressure to please the customer.

Process teams. The model offers a potentially powerful way for Penn to take action at the university level. A few cross-cutting processes will be funded directly and managed across



traditional organizational boundaries. For the moment, Penn will concentrate on two or three high priority processes such as academic innovation and student services. These processes can be considered "institutional bets" with high potential payback. As political implications of the process perspective are worked out, more of Penn's daily life may be organized and funded along process lines.

Costs. Computing is an area of growth and leverage at Penn, one that helps the University save money elsewhere. Penn's campaign to restructure all core administrative activities, for example, and apply the savings to the academic mission, is largely predicated on the capabilities of modern information systems. From a financial perspective, Penn's new model for computing has three aims: targeted investment, cost-effectiveness, and greater control for schools and units. The model tries to unite responsibility and authority where they have grown apart, to reveal real costs where they have become obscured, and to return choice to purchasers where it has been eroded.

The future from here. Consultation with leaders of Penn's units continues. Pilot projects are testing the new model. A steering group is guiding and coordinating the pilots, drawing lessons, and laying the groundwork for transition. ISC and other units are beginning internal restructuring in line with the new model.

Comments invited. Suggestions and comments are invited. Contact Linda May (may@isc.upenn.edu; 215-898-0005). For a fuller version of the model, contact Linda May or see the project's Web site (<http://www.upenn.edu/restruct>).

Appendix II: Campus Pilots Restructuring Computing across Penn

Seven pilot projects are testing aspects of Penn's new model for computing services. The Restructuring Steering Group guides and integrates the pilots.

1. New kinds of learning spaces (*process team*): Begin creating at Penn a range of technology-based "learning spaces" such as classroom/lab hybrids. Build on the success of the Provost's Classroom Committee; seek outside funds.

Team leaders: James O'Donnell (Info. Systems & Computing/Arts & Sciences)
John Smolen (University Life)
Donna Milici (ISC)

2. Support-in-residence for students (*frontline support as close as possible to users*): Pilot the viability of moving to residence-based frontline computing support for undergraduates; lay groundwork for transition. Begin with one residential unit and closely coordinate with broader efforts to restructure student services across the University.

Team leaders: Al Filreis (Arts & Sciences)
Larry Moneta (University Life)

3. Rationalize network central services fee (*network as regulated public utility*): As a first step toward running Penn's network as a regulated public utility, rationalize the "central services fee," which covers network infrastructure. Determine appropriate services, service levels, and fee and explore more strategic approaches to billing. Establish a pilot "public utility commission," or governing board, to make these decisions in consultation with the Penn community.

Team leaders: Ira Winston (Arts & Sciences/Engineering)
Gerry McCartney (Wharton)
Mike Palladino (Information Systems & Computing)

4. Link help desks across campus (*frontline support; services for local providers of computing*): Link help desks across campus by sharing software that tracks problems and solutions. Set common standards and practices for using the software. Arts and Sciences and ISC will initially deploy the software; The Library and the Engineering and Medical Schools will help shape the implementation.

Team leaders: Katie McGee (Arts & Sciences)
Mike Kearney (Information Systems & Computing)

5. Frontline support-for-hire (*frontline support; sell services where it makes sense*): Review and adapt for broader implementation ISC's distributed staffing program, in which local units contract with ISC to locate computing support staff on-site. Explore market needs and develop a formal business plan.

Team leaders: Mike Provost (Veterinary Medicine)
Don Montabana (Information Systems & Computing)

6. Services for local providers of computing (*secondary support*): Begin putting in place a coherent and effective set of services directed to local providers of computing. Determine organizational structures to deliver those services. Define performance measures, service-level agreements, and formal ways to get customer feedback.

Team leaders: Mark Aseltine (Fine Arts)
Mike Kearney (Information Systems & Computing)

7. Facilities management-for-hire (*secondary support; sell services where it makes sense*): Improve, expand, and formalize ISC's program of facilities management, in which ISC provides technical and operational support for computer systems that belong to clients. Treat the recent contract with the School of Dental Medicine as a pilot to learn more about running this service as a business.

Team leader: James Galbally (Dental Medicine)

Appendix III: Project-on-a-Page

Restructuring Computing Services across Penn

Objective. Design a rational, cost-effective structure for organizing, staffing, and funding computing services across Penn.

Process. Four “rapid design” sessions—day-long task force retreats. Look ahead to visions of a desired future, and figure out how to get from where we are to where we want to be. Inform vision with knowledge of current capabilities.

Between sessions, workgroups continue activity in areas of needs analysis, vision, benchmarking, consultation, communication, negotiation, etc.

Sponsors. Provost Stanley Chodorow and Executive Vice President John Fry.

Co-chairs. Professor James O’Donnell (jod@ccat.sas) and Vice Provost for Information Systems and Computing Peter C. Patton (patton@vpisc).

Task force participants. Cross section of Penn faculty, budget officers, computing directors, Information Systems and Computing staff, computing clients, and others.

Workgroup captains. Gerry McCartney and Mike Eley head the Models Workgroup. Ben Goldstein and Donna Milici head the Benchmarking Workgroup. Al Filreis, Ira Winston, and Katie McGee head the Communication and Consultation Workgroup.

Project manager. Linda May (may@isc).

World-Wide Web site. You can locate the project’s World-Wide Web site from Penn’s home page for computing and networking or at URL "<http://www.upenn.edu/restruct>".

