

Planning Input for Oct. 88

Geer planning

Geer Planning input for
Oct. '88 Deadline

to: Directors
from: Dan Geer
date: 15:42 18 Oct 1988
re: Planning input for the 17 October 88 deadline

1. deployment strategy

1.1. questions to be answered

- How do we maximise the impact of our deployment dollars on the market penetration of Athena? Penetration would be measured on two fronts - educational (*i.e.* how to advance education via deployment) and critical mass building (*i.e.* how to ensure that connectivity and coherence are appreciable as worthwhile returns on the Athena investment). (Or, putting it differently, are we a virus or a mate?)
- Is "one stop shopping" a realistic deliverable?
- What are a reasonable number of supported configurations?
- How does direct purchase of workstations apply? This would include servers purchased by others, including services not authored by Athena.
- How to support direct purchased workstations that are of brands other than those directly part of the Athena grants program? (For example, Sun, Apollo and Apple equipment.)
- How might a subsidised purchase program be useful and administratable?
- What role do privately owned workstations and their owners play in the 18 month decision?
- What is the future of the public cluster? A backwater of lowest common denominator machines? A showcase for whatever is the next wave? A place where those resources that cannot be afforded on a mass scale are installed?
- What are the network costs and expectations for performance that we should use for planning? Where does the telephone cable plant fit in?
- What about specialized facilities? Cost sharing arrangements for such facilities?
- What living group deployment is proper and affordable? How does it relate to private purchase and network strategies/expectations?
- Where do research uses and deployment enter in? Do existing users represent a reservoir of deployment potential, or something to be displaced?

- Is it possible to share services with other parts of the campus, *e.g.* file service, and what are the infrastructure requirements for doing so?
- What maintenance strategies go with advancing deployment outside the bounds of the direct grant supported, public cluster?
- What new hardware will our vendors be providing? Must we go the non-disclosure route to find out?
- How will recapitalization be factored into the deployment decision cycle?
- How will deployment be merged into Information Systems?
- What specialized educational equipment should be supported? (For example, projection equipment.)
- Do adequate reservoirs of public space exist to expand? Or do public spaces become obsolete within the life of this plan?

1.2. methods for answering

- Define the tough limits on the systems first, then work up as many scenarios as we can. This will require new, possibly non-disclosed information from our vendors.
- Gedanken experiments about the nature of a service environment with highly distributed control *vs.* highly centralized control. The differences in outcome for deployment between the two results should be indicative of a planning direction. This is not to say that the two poles, highly centralized *vs.* highly distributed, are the only possibilities. Rather, the point is that the solution space lies between the solutions that the extremes these two models define.
- Retest the hypothesis that student purchase is desirable. Are the grounds purely capitalization / financial?
- Is deeding workstations to departments, *etc.*, a way to build an investment-driven consensus for the 18 month timescale? If so, identify likely sites and favor them.
- Determine if we care about avoiding price competition with other on-campus suppliers.

1.3. who will be responsible for solution?

1.4. when will solution be completed?

2. end-user interface problems

2.1. questions to be answered

- Can anything real be done in 18 months? Can the problem even be succinctly defined in that time scale?
- What are the real needs?
- What priorities should be attached to meeting the individual needs and how can these priorities be set so as to support the right decision in 18 months.
- Is a one-size-fits-all strategy a rational possibility? Assuming that is an absolute impossibility, answer the question “Who is the customer for this?”
- Is there an interface that promotes educational use, *per se*?
- Is waiting for the {OSF,...} a rational option?
- Who or what are the end users; what are their characteristics?
- What do you want to interface to? Is it text, files, the operating system, widgets, databases, *etc.*, *ad infinitum*?

2.2. methods for answering

- Attempt to influence any standard that the OSF brings forth, or that appears to come from some other important standards setting / promulgating body.
- Debrief inside Athena and faculty developers, paying particular attention to shortcomings that were determinative of the outcome of the individual project.
- Do a survey of other university solutions. Evaluate them on a common, predefined basis.

2.3. who will be responsible for solution?

2.4. when will solution be completed?

3. application development tools for faculty & developers

3.1. questions to be answered

- What are the real needs?
- Should we optimise for long-term quality development or throwaway apps?
- Can tools help produce more applications?
- Is the purpose of tools to maximise the performance of the best designers and developers or are is the purpose to

make use of otherwise mediocre talent?

- What applications have the highest educational promise and what tools correspond to meeting their particular promise?
- Is it rational to assume that tools can substitute for skill and sophistication?
- Do faculty and students have the same needs? In other words, who is the customer?
- Are future faculty or current faculty the more productive focus? (This is the argument that it is today's end-stage graduate students that are the most cost-effective educational innovators, and the best target for innovation by Athena.)
- What is the relation of these posited tools to export and import of educational software?
- How does one define tools that maximise self-sufficiency of their users?

3.2. methods for answering

- Industry surveys, especially of computer-aided software engineering (CASE) and other tools that purport to lead to lasting success. Or are these tools too sophisticated for the faculty developer class?
- Ask (and answer) the priority question, *i.e.* are there a few applications that it is worth supporting directly - if so, determine some method of defining them.
- Do a careful, structured debriefing of the application developers both in-house and among the faculty. This may suffer from small sample problems...

3.3. who will be responsible for solution?

3.4. when will solution be completed?

4. tools & support costs for large scale deployment

4.1. questions to be answered

- Is ten-thousand the correct design point? (In other words, is full deployment at M.I.T. a sufficient aim, or cross institutional usage? Is this the time for a Boston-area educational computing model?)
- What degree of distributed control is a requirement?
- Will network technology impact our sense of scale to bandwidth ratio, and how?

- What standards will arise and to what degree will they require reengineering the Athena hardware plant? (This is sustainability at scale that is being questioned.)
- Are workstation and server configurations up to the task?
- What is the role of vendor software and our influence on it?
- Are proportionality constants currently attained sufficient? If not, what proportionality constants are needed? (For example, we are now two orders of magnitude off the one-wizard-per-cpu rule of thumb of traditional Unix management. Do we need to be three orders off that rule?)
- Where does independent purchase and self management fit in with respect to scalability of the environment?
- Do we need different operational models to make this work at higher scale? (We currently overwork a single hyperwizard, underwork others.)

4.2. methods for answering

- Push deployment and see what breaks.
- Export our software aggressively to other large scale sites, and see what breaks.
- Get better feedback from our own experiences.
- Work particularly closely with our sponsors and with the OSF on the scalability considerations of AIX.
- Participate in the Nationwide Filesystem experiment.
- Define a service agreement with other deployment sites, such as EE/CS, and determine the operability and maintainability consequences of that agreement. Develop accordingly.
- Reengineer the Unix workstation based on what we now know.
- Do all the development now known to be required for the 10K scale. (For example, dynamic network address assignment and network error logging.)

4.3. who will be responsible for solution?

- Athena Systems Development holds the responsibility for the workstation software.
- Athena Operations holds the responsibility for determining affordability and functionality constraints on long-term operational models.
- Athena User Services holds the responsibility for determining the practical maximal skill level at deployment sites that User Services can support.

- MIT Information Systems holds the responsibility to determine the workstation configurations receiving official blessing to be ones that are consistent with the computing model which I.S. will encourage users to adopt.

4.4. when will solution be completed?

5. information management

5.1. questions to be answered

- What volume of pent-up demand for massive databases now exists? What part of that demand is generic to the educational process?
- Will wide-usage databases evolve at rates consistent with yearly, read-only distribution methods, or are active read-write access methods important?
- Are these databases to be offered in conjunction with the libraries or in opposition to them?
- Are analytic tools commensurate with the scope of these databases available on the MIT campus? Are they affordable and accountable?
- Where do traditional suppliers of statistical and retrieval services stand on the distributed computing issues?
- If non-Athena suppliers will be part of the service environment, how do we capitalize their needs for self management, tools, and up front costs? (Costs broadly defined.)
- How is cost recovery to be managed, including cost recovery for externally provided information services (such as Medline)?
- Is our inter-user communication model up to the job? (Treating inter-user communication here in the broadest sense, the sum of Zephyr, email, OLC, discuss, *etc.*. These are all parts of general information management.)
- What browsing and/or navigation tools are appropriate? What cognitive models do we wish to support, and does local expertise exist to guide Athena?
- How does direct capture fit into the information management model? (For example, the lab workstations.)
- What is the optimal on-line help system?
- How do active filters, such as the Information Lens, apply to the information management needs that real Athena education uses require?
- What is the future of printed documentation, and how important is it to the future of the project?

5.2. methods for answering

- Coordinate views with the traditional suppliers, *e.g.* Information Systems.
- Determine if the lack of these services is a clear and present problem, or whether it represents unrealized opportunity costs. Debrief faculty. Determine if actions being taken at other universities have bearing.
- Define a set of software services that could be thought of as minimal - and determine if those can be scaled to meet demand.

5.3. who will be responsible for solution?

5.4. when will solution be completed?

6. interconnect to other systems

6.1. questions to be answered

- How far down? How far up?
- What is the educational aim? Information sharing? As a crutch for deployment? Political defense?
- Can the ISDN cable plant be pressed into service for connectivity? What lead time and cost issues apply?
- Will the PC and Mac type machines fade away or not? Will they be adequately represented by compatibility modes on real workstations?
- Is there much educational value to faculty working at home?
- Do students in apartment living arrangements represent an unmet fairness need with respect to other deployment plans in organized living groups and on campus?
- How do we bound the expectations that connections will bring? How, particularly, do we avoid unintended devaluation of our model of computation?
- Does Athena have a role as a "change agent" for M.I.T. as a whole? Does that include encouraging upscaling existing low-end machine users?
- Does extension to other environments undercut coherence? Assuming the answer must be yes, what is the cost to benefit of that step?
- What does easy interconnection mean to issues of import and export of software?
- Is there vendor interest, especially with respect to our sponsors, in interconnect to other classes of machines? Are

they taking active steps to obsolete existing computational systems?

- Supercomputer access can be as simple as providing an X client. Is this sufficient? What are the bandwidth issues? Will Ethernet speeds be too much of a bottleneck to permit real use of supercomputer resources? (“Real” as defined by the client base.)
- What are the real effects of differential bandwidth connections on patterns of use? Does that matter educationally?

6.2. methods for answering

- Get reliable figures for living arrangements and for existing ownership. Perform a proper survey to separate utilization of existing investment issues from situations where the purchaser would repeat that purchase one year from the present.
- Retest the oft-stated resolve to require student purchase. Reevaluate the lease of current model machinery rather than supporting obsolete or non-compliant machinery.
- Interview all other providers of basic computing services on campus to see if someone else has a better solution for solving this problem.

6.3. who will be responsible for solution?

6.4. when will solution be completed?

7. publications, *et al.*

7.1. questions to be answered

- What are the aims, specifically, for public relations?
- What is the audience?
- Are those aims separable? For example, publicity for getting the right answer in 18 months vs. publicity for getting industry acceptance of Kerberos. How many kinds of right answer are there?
- Is it appropriate for publicity to be funded on a per group basis or through the Directors? (For example, travel costs to conferences where Athena presents.)
- Do the supporting vendors have a particular role to play? Should we solicit or accept their help?
- What opinion leaders on campus can be influenced? Are department heads among those?

- What value can be derived from those whose Athena projects succeed? Those who fail?
- How important is quality control with respect to public relations? Where is the point of diminishing return and how do we detect it?
- What is the effect of response time on public viewpoints of Athena? (Response time all the way from key stroke latency to the time it takes Athena to respond to a user communication of any form.)
- What is the effect of response time with respect planning? (In other words, what is the effect of the delay time from planning to application to service on public perception?)
- What is the role of Customer Service? Should we have a department specifically for customer service? Is that what Operations and User Service are? Why?
- Is OLC service at a sufficient level?
- Is hotline service at a sufficient level?
- What data collection changes can better support public perception of Athena? (The present information amounts to statistical malpractice.)
- How do we do benefit accounting on an equal strength to cost accounting? How do we define the beneficiaries, and do they map one for one to the audience of the public relations focus?
- Why are there no efforts by Athena to show itself in general educational computing venues? Do we have any reportable successes in this area? (Reportable meaning potential for papers that would pass conference and journal referees.)

7.2. methods for answering

- Assess the degree of penetration we have for each product, the importance of an industry acceptance of our way of doing things, *etc.*, and target paper writing accordingly.
- For opinion leaders, determine if they have single issue foci and whether this is addressable.

7.3. who will be responsible for solution?

7.4. when will solution be completed?

8. graduate students

8.1. questions to be answered

- Why is this a separate category?
- What are the real needs?
- Are graduate students to be considered educational or research users and does this matter? (Licenses, for example.)
- Will graduate students usage be a straight line percentage increase over undergraduate uses? Exponential? How comparable are they? Does data exist, here or abroad?
- Are desks of grad students a deployment frontier? If not, how much more public workstation resource is required?
- Are there fairness issues with respect to varying levels of departmental budgets?
- Where does cost recovery come in, and is that different for this class of users?
- Where do traditional computing providers come in, and is that different for this class of users?
- Do graduate students represent a group of sufficiently motivated (or motivatable) individuals to permit sharing the operation load? Is the public model a sufficient workstation model for this group?
- Does a modern graduate education now presume familiarity with computing? What percentage of dissertations involve an important computing component? Who now bears the teaching load for bringing the graduate students up to speed?
- Are portions of the management of the graduate programs better done with computing resources? (Schedule setting, submission of important milestone reports, communication with the Committee, *etc.*)
- What disparities now exist between graduate students of different departments? Is it Athena's role to preserve, lessen, or exacerbate those differences?
- Do graduate students have a special interest in the information management issues discussed elsewhere?
- How can Athena provide computer supported cooperative work (CSCW) and is that an appropriate component of any approach to graduate students? (Arguing that the nature of laboratory science is cooperative.)

8.2. methods for answering

- Figure out how we can make do without scientifically valid sampling, which can not afford time-wise.
- Determine the competing resource costs to do this versus some other signal use of Athena's resources, living group

deployment, say.

- Consult with Department leaders and draw them in, somehow.

8.3. who will be responsible for solution?

8.4. when will solution be completed?

9. strategic courseware dev

9.1. questions to be answered

- What did we learn from doing this once? How solid are the inferential bases for what we learned?
- Does specially supporting particular courseware make it better? Better in proportion to its increased cost? (Define better to include educationally effective, more durable, utilising a greater fraction of the power of the Athena model, sophisticated and intuitive at the same time, and contributory to the right decision in 18 months - and reanswer the question.)
- Is it important to support courseware driven by tenured faculty? (Compared to the impacts noted to date.) Or is it better to support the next wave faculty (end-stage graduate students)?
- Would it be more cost effective to support the development process than particular development? Would dollars spent on a Design and Evaluation Guide for Quality Educational Software be more productive? Would we know what to say in such a book?
- How do we separate wants from needs?
- Does any benefit besides resource conservation accrue from the Stanford model? Is this a provable assertion?
- Who, exactly, on campus is the arbiter of educational expertise? How can Athena play or coopt that role?
- Are we, in fact, still committed to doing courseware? If so, what are our actual credentials and what exactly is provably correct to do? How do we proceed?
- Should Athena provide design and/or code review to faculty projects?

9.2. methods for answering

- It's time for a *post hoc* analysis of what we learned in the curriculum development effort of the past five years. This would include infant mortality figures, characteristics of success vs. failure-to-thrive, and whether projects

can be supported in the long term.

- Determine, if possible, the list of three most promising but unfunded courseware development projects. Are they so good as to be worth putting all our eggs into their baskets?
- Set up a design principles guide that determines (in advance) a development strategy that results in Athena maintainable courseware.

9.3. who will be responsible for solution?

9.4. when will solution be completed?