

A Lifetime Burning in Every Moment

By Mark A. Luker

AN INTERVIEW WITH **KENNETH J. KLINGENSTEIN**

Kenneth J. Klingenstein, Chief Technologist at the University of Colorado, Boulder, and Project Director of the Internet2 Middleware Initiative, is the recipient of the 2003 EDUCAUSE Award for Leadership in Information Technologies. Klingenstein has been involved in computing technology for almost three decades. One of a handful of leaders who have put and kept academia at the forefront of Internet development, he is widely respected for his breadth and depth of knowledge, his uniquely accurate visionary abilities, and his skill in identifying and communicating the salient problems and issues embedded in complex, often chaotic advanced network environments.

One of the true Internet pioneers in the higher education IT community, Klingenstein was involved in early NSFNET development, chaired the Federal Networking Council Advisory Committee, and has participated in many federal advisory groups on network policy and technology. He served on the board of the Corporation for Research and Educational Networking (CREN) and was among the key initiators of Public Key Infrastructure (PKI) security activities at a national level. He has served on the governing bodies of CAUSE, the Common Solutions Group, and the Coalition for Networked Information (CNI), as well as of regional and national networks. He was one of the first people in the United States to recognize the importance of community and school networking: he received the first NSF grant issued to systematically network a school district and create professional development opportunities for teachers; and he received a similar community network grant, from the National Telecommunications and Information Administration, for the Boulder Community Network. His many workshops, talks, and publications convey the intelligence, authority, wit, and clarity with which he perceives the complex IT environment.

Mark A. Luker, Vice-President of EDUCAUSE, recently talked with Ken about the early days of his career, some of his past and current activities, and his thoughts on the higher education IT issues of today and tomorrow.



Luker: Ken, in advanced networking circles, you are known nationally, even internationally, as “Mr. Middleware.” How did you get your start in information technology for higher education?

Klingenstein: Mark, you and I were trained in computer science even though we received our Ph.D.’s from the Math Department at the University of California, Berkeley. When I was a professor at the Colorado Springs campus of the University of Colorado, I became chair of a faculty committee that was evaluating computing services. The committee reached the unanimous opinion that there needed to be a change of leadership, and when I delivered that report to those representing the governing authority, they acted on it. Then they turned to me and said, “Well, you brought this in, you sweep up.”



Luker: That could explain why you have such a good appreciation for the faculty issues of teaching and research.

Klingenstein: As well as a mastery of the passive tense, which we used a lot in faculty committees.

Luker: What were the main issues that you faced in information technology in your first IT management job?

Klingenstein: I was at a smaller campus, and so there was the overarching issue of the smaller campuses vs. the flagship campus. Later on, I wound up being the information officer at the flagship University of Colorado campus and hoped that I would remember the perspective from the lower campus. There was the issue of building localized independent computing facilities on campus, and there was the always interesting question of where academic computing and administrative computing overlapped. We had a centralized, multi-campus administrative computing shop, and I remember that the first time I donned kneepads was to convince the management of the multi-campus shop that e-mail was an academic utility and a valuable one. I very vividly remember the response was that no one would ever use a service like that.

Luker: Speaking of e-mail, how did you get into networking? I understand that you were active in the early development of NSFNET.

Klingenstein: Actually, that goes back to my faculty days. There was the need to teach a course in networking. The department had never taught one before, and as the most-junior faculty member on the staff, I was given this worthless and throw-away topic of networking and told to teach it. A couple of years of teaching that subject gave me at least some academic background, and then I became the CIO, so I was able to act on that background. Then when I moved to the Boulder campus, Boulder was one of the charter members of the John von Neumann Supercomputing Center (may it rest in peace) at Princeton, and so suddenly there were the issues of running networking capabilities from Princeton back to

the Boulder campus and of building local campus infrastructure to disseminate that supercomputing network access. In my first year at Boulder, I spent a lot of time learning about satellite dishes and getting the satellite to JvNC up at just around the time that NSFNET was providing an alternative and superior path.

Luker: How did you get started working with the National Science Foundation (NSF)?

Klingenstein: Like a lot of things in my life, it was luck. I got invited to give a talk at the Kennedy School of Government in the early days of networking, and one of the attendees was Chuck Brownstein, who at that time was NSF deputy assistant director for Computer and Information Science and Engineering (CISE). He heard me, and I guess he smelled fresh meat and passion.

Luker: What role did the Federal Networking Council (FNC) Advisory Committee play?

Klingenstein: What we were trying to do there was persuade the federal agencies—which at that time were running lots of independent backbones to places close to campuses and then asking the campuses to pick up the delivery of that independent backbone to individual desktops—to provide a uniform national backbone, to base it on TCP/IP, and to have that be the generalized research vehicle. Then we would just plumb that single backbone widely across campus.

Luker: I take it that the FNC, in its way, and the NSF program officers were what we would today call Internet governance—or the closest thing to it.

Klingenstein: It was certainly the federal backbone network governance—“governance” in that loose cooperating sense among federal agencies. The advisory committee brought a beyond-the-beltway perspective to the mix. It was a clash of views, but there were remarkable people on both sides, and the results speak for themselves.

Luker: In recent years, there has been a

concerted push to wire every K-12 school—and, some argue, every classroom—in the United States. I understand that you got into that discussion very early in Colorado.

Klingenstein: The Boulder Valley School District was the first school district to receive an NSF grant not only to do systemic district networking but also, and even more importantly, to do teacher training and evaluation. It struck me as significant because when NSF awarded us the grant, they doubled the amount of money that we had asked for. They said, “We want you to do the job right.” In retrospect, NSF was looking for this kind of

“From the very beginning, we’ve had one compelling reason to network, and that reason persists: the strong collaborative forces within higher education.”

systemic, physical, and professional educational networking opportunity. For me, living in Boulder—which at that time was the most-connected city in the country measured by, I think, dial-up ISP accounts—the opportunity felt like an obligation. I didn’t think, “Oh, what fun!” It seemed that if we were going to try to implement networking in K-12, the place that it had the best shot to succeed was Boulder, and so out of obligation more than spare time, we put together the proposal. Then NSF came back and said, “We like the proposal, but we think you can do more, and here’s what we want you to do.”

Luker: Back then, in the days of CSNET and NSFNET, the research and education community played a leading role in both the design and the deployment of the Internet. Since that time, we’ve seen the explosive growth of the Internet industry, a growth that has left higher education with a much smaller share of the investment and overall activity. How has this changed our role?

Klingenstein: From the very beginning, we’ve had one compelling reason to network, and that reason persists: the strong collaborative forces within higher education. Now that the desire and climate for collaboration has been addressed at the network layer, it has percolated up to the middleware layer, and so much of the work that we’re doing here is around tools for inter-institutional collaboration. When we talk about federation to people involved in the major efforts in industry, they always say, “Well, we are pleased that this is coming out of higher ed, but we are not surprised because we do understand how much you people need to collaborate between institutions. We have the same needs that you have, but for you they’re more urgent.” So I think leadership in networking by higher education persists, it’s just higher up the protocol stack than it was a few years ago.

Luker: Some people feel that higher education almost declared networking victory several years ago after we achieved standardized, efficient ways to support high-speed networks throughout our campuses and between our institutions. In hindsight, we see that we had not realized what other components were needed. What were we missing?

Klingenstein: I think there are two broad categories. Certainly, we were missing the middleware pieces, the infrastructure to enable collaborative application. And then the other thing that we let go of too early—and that we’re now trying to get our hands around again—is the premise that the network needs to carry not only the applications of today but also the applications of tomorrow. So many of our efforts to secure our networks today are utterly necessary activities, but at the same time, they’re doomed to failure because either

(1) they will fail due to newer innovation or (2) they will succeed and the Internet that results will fail in its mission to permit innovation. For example, firewalls provide immediate network security benefits but will soon be undermined by encryption. At the same time, to the degree that firewalls are successful, they can block new next-generation applications like desktop video and grid computing. If we keep to the current security approaches, no new next-generation applications will come to the fore because they won't be permitted on today's Internet.

So I think we have to pick up the gauntlet in at least two areas. One is in the middleware space, where I think higher ed has responded. The second is in network security, where we need to combine new requirements for security with the persistent goal of allowing invention and innovation to happen on the Internet. We need to pick up that leadership and show people that we can achieve both goals—network security and network innovation—at the same time.

Luker: In 1999, you were loaned from the University of Colorado at Boulder to Internet2 to lead a concerted effort to establish middleware standards and products to support a broad range of networked applications in higher education. What is this effort about?

Klingenstein: On a technology level, it's about creating enterprise authentication services, enterprise directory services, and in the next year or two, enterprise authorization services that can also be used to support scientific groups that do team science. Campuses need to deploy a limited number of core software systems that work with each other, that work with the lower layers of the protocol stack at the network layer, that work with the upper layers at the applications layer, and that work consistently between campuses to encourage the collaboration. So, one whole piece is to install these software systems.

One thing that's different, however, from working at the network layer is that these software systems are capturing more information about people than about machines, and so there are huge policy issues that are much more domi-

nant at the middleware layer. So, a fair bit of the work is about encouraging campuses to make these investments, persuading the owners of data silos to open those silos to other applications, and putting in place privacy and security policies for all of the individual attributes that we're suggesting campuses keep online.

Luker: In the past, the word *enterprise* has often been used to refer to overall standards and designs for business systems and for the administrative system of an institution—as in ERP (enterprise resource planning). Can we think of *enterprise*, in your sense of the word, as extending that kind of planning to support the academic functions as well?

Klingenstein: Certainly. And it's reflective of the fact that there are more and more systems that cross the academic/administrative boundary. A course management system, which is certainly academic, has huge amounts of administrative data being fed into it and needs to run

with the professionalism of an administrative application. Portals are often academic in nature but need to be operated with the solidity of the administrative shop. It is the crossover applications—the ones that are hard to differentiate between administrative and academic—that suggest that the word *enterprise* now needs to span both sectors of the community.

Luker: You have brought together a thriving community of international talent in the middleware project. How do you practice leadership in such an environment? Why do the others participate?

Klingenstein: I'm not sure I know the answer to that one. I know that part of the reason they participate is because the territory is so rich and fertile that accomplishments are attainable. The past success motivates the future energy. I often think about middleware as this super-saturated solution; any string placed into that solution would have precipitated wonderful crystals, such as what we are

seeing emerge today. That said, people have told me that there is something that I've added to that mix. It may simply be that I'm not good enough to be a technical expert and not good enough to be a politician but that I am half-good as a politician and half-good as a technical expert—and that's what is needed for leadership.

Luker: Can you tell where we're going in middleware? If not, how do you know where to lead? What are you doing to figure out directions and emphases?

Klingenstein: The middleware effort has been led by the members of a steering group called MACE (Middleware Architecture Committee for Education), and all the technical direction comes from their expertise. These are the best and the brightest of the campus middleware architects, so I really need to defer to them for much of the vision. I should point out that middleware will ultimately plateau, in the same way that most other technology areas have plateaued. People don't talk with great excitement about TCP/IP anymore. I suspect that in a couple of years, directories and authorization will be as pedestrian as TCP/IP is today. I don't know that any area of innovation goes on forever. I think that we'll reach a plateau, but I think we'll know we've hit the plateau when the MACE members say, "OK, we're moving on to other kinds of issues."

Luker: And sometime after that, the rest of the world will forget that they didn't always have enterprise directories.

Klingenstein: That's an excellent point, Mark. When you're trying to sell transparency—and middleware is a fair bit around transparency—people see that transparency exactly when they don't have to go through a door or a barrier or, in the case of middleware, an authentication process. Suddenly, they get to something they wanted to go to without the intervening hassle. They'll appreciate that ever so briefly, and then it's going to be, "What have you done for us lately?"

Luker: Ken, you're now one of the principal investigators in the NSF middleware

initiative. What does NSF bring to the equation?

Klingenstein: NSF has brought some really important funding, but even more significantly, it has brought a needed imprimatur. If we were promoting standards and recommendations by saying, "Fifteen bright people on campuses have said this is a good thing," we would get, I would guess, somewhat limited traction. With that NSF label, we have achieved extraordinary traction, a wide variety of campus involvement, and also a much bigger international presence than we ever expected. In the original proposal we had very modest goals, especially

“I suspect that in a couple of years, directories and authorization will be as pedestrian as TCP/IP is today. I don't know that any area of innovation goes on forever.”

internationally. I think the NSF imprimatur has given us much more standing internationally, and that's been very helpful.

Luker: You have recently been asked to help Internet2 mount a major campaign involving issues of network security. How will you approach that assignment? Which issues are most important in the context of advanced networking?

Klingenstein: Let's talk about both process and product. In terms of process, I think it will be similar to what we've tried in middleware: gather together the best and the brightest from the cam-

puses. We're doing this right now to form an advisory group that would prioritize and make recommendations on direction and perhaps even point out expertise that could be tapped for the work. As for the product, what we're trying to do is complement the excellent work already being done by the EDUCAUSE/Internet2 task force in awareness and education, in policy, and in current practices. But higher ed has not given proper attention to the issues of preserving innovative capabilities on the Internet, of understanding how security should be deployed in ways that permit scientists to still have high-bandwidth or low-latency interaction globally, or of developing some new testing tools for denial-of-service attacks and understanding how they can be deployed on campuses or between campuses to provide early alerts when a DDoS (distributed denial-of-service) attack is occurring. The higher ed community, because of its high-bandwidth pipes, has a particular obligation to make sure that those pipes are secured, especially when it comes to denial-of-service attacks. And so I think we have a fair bit of work to do in learning how to diagnosis DDoS attacks earlier and to implement responses.

One last area of work involves the relationship between middleware and network security. There is a little bit of leverage that we can gain by using network authentication and authorization—either of devices or of individuals using devices—to help improve the security environment. One example might be role-based, policy-oriented personal firewalls. So many of us in the higher ed community believe that the best defense, given how soft our perimeters are and how fungible they need to be in the face of the academic requirements, is to put as much security as possible on the desktops. But that security can't be uniform. It has to be commensurate with the role of the individual. If the individual is a researcher using a grid, the researcher is going to need some ports open that students and staff may not need open. So there are some very interesting opportunities in using the roles that we're oriented toward in middleware to shape the configuration of desktop firewalls in ways that are transparent to the end user

but that provide much greater security. I look for efforts in those kinds of symbiotic arenas that join middleware and network security.

Luker: This could provide both the needed flexibility and an efficient, reliable way to make the correct decisions without an army of system administrators tweaking every machine.

Klingenstein: Precisely.

Luker: How is security related to privacy? You're a well-known advocate for privacy.

Klingenstein: That's a complex relationship. At a conceptual level, security is about establishing the identity and key attributes of a user, and privacy is about protecting that information from inappropriate access by others. At an operational level, security and privacy are often seen as two sides of the same coin; thus we need to make sure that our security efforts do not needlessly impinge on privacy. There has been some interesting dialogue recently about NetFlow statistics and how public or private those statistics might be. We're starting to see the "personally identifiable information" phrase applied not only in the typical way that colleges and universities think about it—in terms of either academic management or health care—but also in traffic patterns on the Internet and the rules that should govern that traffic. How can we keep those logs in ways that preserve our ability to detect problems in host devices on the network but at the same time not needlessly compromise the privacy of the individual who might be running that host? This involves a very difficult set of issues.

Another operational place where this security/privacy dynamic has come forward is in middleware diagnostics. Most of these middleware tools are privacy-preserving, but the log files that they generate totally undermine that privacy. So when we go into a diagnostic mode and want to harvest the number of different log files to understand why a middleware transaction didn't succeed—that is, why a user wasn't able to access a desired resource—we're going to have to be very careful that the forensic activity does not needlessly impinge on privacy. I don't

think we know the solution to this problem, but at least we've formed some succinct expressions of the question. It's often good to have targeted use cases to help motivate solutions, and I think we're to the point that we at least have the good use cases.

Luker: On the other side, of course, systems without security can't guarantee privacy of their information.

Klingenstein: Right. That's an excellent point. We see tremendous compromise of privacy because the system itself is insecure.

Luker: The subject of collaboration tools was mentioned by you earlier, and it's mentioned nearly every time someone discusses advanced applications. Your leadership has been key to a number of successful collaborations. What

collaboration tools do you use today with MACE and the other groups?

Klingenstein: Because we don't have the full middleware structure in place, we're still pretty primitive in this area. We do a distressing number of conference calls, and we use listservs even more often. We are now fairly close to point-and-click desktop video-conferencing, and we have some enhanced Web-authoring tools coming down the road, but right now we're too busy inventing the tools to spend much time using them.

Luker: What advice do you offer aspiring leaders in information technology in higher education?

Klingenstein: To maintain their passion. To be well versed on the fundamentals. There's not much new under the sun. I wouldn't suggest we go back and get punch-cards and try to figure out how to use them, but I don't think the frameworks and grounding of computing have

changed much. For example, I do a lot of work with the grid, and I see the grid come up against interesting issues that were first explored in operating systems twenty-five years ago—issues about deadlock and the deadly embrace. It seems that as we move into more complex systems, the simple problems that plagued simple systems also plague complex systems.

Given that we are dealing now with a far more complicated environment, we all need to remember to keep things as simple as possible. Complexity kills. In Shibboleth, a collaboration tool that higher ed has recently developed, we put in a tremendous number of knobs and controls that both system administrators and end users can turn in order to get the tool to work exactly as they want. Yet I think we should spend years growing into that kind of fine-tuning. Even when technologies present themselves with a rich and complicated set of options, we need to keep things extremely simple and grow into those options over time.

Luker: One final question, Ken: Are you going to use this award as an opportunity to relax, to take your foot off the pedal?

Klingenstein: Is it true that you get only one lifetime in a lifetime? If so, maybe I will. But there's a wonderful quote from T. S. Eliot about how there is "a lifetime burning in every moment." I see more of those moments still stretching in front of us. *e*

RELATED RESOURCE



The EDUCAUSE/Internet2 Computer and Network

Security Task Force is working to increase awareness of IT security issues across higher education (<http://www.educause.edu/security/task-force.asp>). The absence of common middleware solutions is being addressed through the work of NMI-EDIT (<http://www.nmi-edit.org>).