Application-Driven Technology

Ted Hanss
University Corporation for Advanced Internet Development
Ted@internet2.edu

Thank you. What you see on the screen here is the URL apps.internet2.edu/talks that I will repeat on my last slide. So if you want a copy of my slides, including the pictures and the movie clips that are included in the presentation, you can go and grab your own version. Again, I ll repeat this on the very last slide if you don t have your pens out right now.

Basically, I am a cheerleader for the kinds of activities that are going on inside the universities. There really isn t a central activity for Internet2 unless you think about all the participants and all the universities working together across our community. And the excitement I derive from working on this job is from getting to meet and work with the faculty. What I want to do in the next few minutes is just outline for you some of the development efforts going on inside the member institutions.

One of the first questions I get is: What is an Internet2 application? And the way I answer it is by pointing to those applications that make a qualitative and quantitative difference in how we do teaching and learning and research and clinical applications on our campuses. It s a you know it when you see it environment. Except the one distinction is that these applications don t run on the commodity Internet. They require an advanced infrastructure in order to run. And that s what distinguishes them.

Now, the history of high-performance computing, as you've heard from some of the history examples in the timeline shown, is in the attempt to tie supercomputer centers together. But the Internet2 project is about all that we do in higher education. So, yes, the science environment has been the leading area for these kinds of applications. As you'll see from some of the demos out in the lobby, arts and humanities applications also play a strong role in this environment -- whether you re doing it, again, for instructional in the classroom, for research purposes, and so forth.

Another question I get asked is What s the killer app for Internet2? Going back to involvement in the development of the Internet, people couldn t predict the advent of the web in a certain time frame and how that led to the proliferation of the applications that many of us use today. But when I think about the killer applications for Internet2, I think about them in terms of, What are the killer attributes for our environment? And the first one I d like to talk about is interactivity. That is, using the ability to interact with other people in a collaborative context with access to common resources like databases, remote instruments, and so forth. So it s a rich, interactive, collaborative environment, unlike the click and wait environment we have on the Internet today. We also continue to have large-scale computation. But instead of being centered in physical facilities and machine rooms, it really is a network-centered way of tying together these resources. And as you heard in the previous presentation by John, the whole area of shared virtual reality or teleimmersion, is the new way of looking at our data. I ll be touching on that in some of the examples. So it s a combination of collaboration, interactivity, common access to resources, uniting resources across the network, and looking at data in new ways, that make up the killer environment.

Now one of the things I have mentioned specifically in the arts context is that I think the killer app for the artist is high fidelity. That is, the ability of Internet2 to communicate content without losing what is behind the delivery and the creation of that content. So the integrity of the artist's message, the vision, the communication, is what can be maintained across the kind of technology.

And so again, as you can see from some of these examples, high fidelity environments for making those communications possible. So let me go through snapshots of some examples to give you a flavor for the kind of things that are being developed by our member institutions.

Coming from Indiana University, which has the largest School of Music in the country, they ve had a challenge of making available resources for their students, allowing them to sit and listen to various musical clips. They went through a model of handing out cassette tapes and then transitioned to digitizing the music but only making it available from desktops within a single room in the Indiana University Music Library. What we did with Indiana was to take those resources and put them within the Internet2 infrastructure so that anybody could have access to it across the network. Now the distinction between this and how we ve done streaming audio on the Internet today is that music students don t want to wait 30 or 45 seconds for the music to buffer up on their local machine in order to play it. They want to jump around and listen to clips, compare clips, with the same kind of response you get from a CD player. And so the important thing about this application is not its high performance (it s only about 400 kilobits per second) but it's high interactivity. The students, again, can, from wherever they re sitting, jump around in the music archive and get about a one-second response like they could get if they can just jump around from various clips inside a CD player. So, again, stressing the importance here of interactivity, not necessarily the exclusive benefit of high bandwidth.

Now, an example of outreach, in terms of bringing this communication to broader audiences, comes from Gallaudet University, the school for the deaf and hearing-impaired in Washington, D.C. There they ve long recorded lectures and made them available as a lending library both on campus and other institutions that have deaf and hearing-impaired students. What Gallaudet is working now on, as part of an Internet2 project, is digitizing these video clips and making them available on demand both to the campus and beyond as learning resource tools for students. And here the high quality video is necessary to carry the facial expressions and the various nuances of American Sign Language. The little postage stamp video in the corner is not suitable for carrying that information. You need full motion, nearly full screen video. And that s the kind of thing that Internet2 enables.

One of the applications that will be demonstrated while you re here is the Space Physics and Aeronautic Collaboratory that supports the groups upper atmospheric scientists who traditionally traveled in the foreground in these trailers to look at instrument data gathered from the ground base radar and from satellite imagery. And I like to use this as an example of an application where the network has made it possible for the scientists to collaborate with each other without having to travel up to Greenland. But instead where the scientists and instruments around the world can actually interact with each other via a web browser interface. Where they can have access to this instrument data, both their own instruments and the instruments of other scientists, they can chat and collaborate with each other. This is all recorded for later playback and viewing.

Well, I think there are a couple of interesting lessons that have come from this. I ve heard the developers talk about the fact that it is no longer just a select group of folks who were able to travel to Greenland to participate in the science. Actually the majority of users are the graduate students who are now able to engage in scientific work virtually shoulder-to-shoulder with the leading researchers in their field. These students are gaining the experience of doing science much earlier in their careers. And so an interesting second order effect here is not just making the virtual laboratory for the already involved scientists, but expanding the access to it.

Now another thing we ll talk about too is also the impact of how this data can be visualized in new modes when I talk about teleimmersion. Now some of the challenges you think about when we start making this available: we talk about larger images being carried more easily across the network. As an example, MIT and many schools are working on geographic information systems, bringing it down to finer and finer resolution. And looking at images from both satellite imagery and from airplanes and mapping these together so you can fly through the various spaces and look at information.

And then one of the things MIT is working on is an overlay in it with other kinds of detailed statistical information about neighborhoods. For example, retail environments or where people live, whether low income or high income. And you see some very interesting things. One of the first things I did here was similar to web browsing. The first thing you often do when browsing is to sit down and type your name in to see how often you re referenced on the Internet. Here I went to search the place where I lived when I went to school in Boston, and I found my apartment here in Boston. Then you start realizing, as you re moving forward with this is what kind of implications there are for this.

Maybe this is something with far-reaching implications.

As we start getting down to environments of one meter resolution for images, can you imagine what will happens when you can move information like these high-resolution images, around at the speed of light. You start getting phone calls from someone who says, I noticed that your lawn looks to be struggling right now, can I sell you fertilizer services? Or, There s heat escaping from your roof, is it time now for a new roof and insulation? And this is the kind of information that will be at people s fingertips with the kind of facilities we re putting together and support both the capture and the distribution of such content. And for some maybe it s not as interesting for people to be able to find where I live, when they also mirror it with other kinds of data resources out there.

Now, again, I echo what John Evans said about being able to get to the scientific instruments across the network. Some very compelling applications, including one here up on the north end of campus is a scanning electron microscope that s buried in the bowels of the building across the grass there. And this one is a fairly affordable scanning electron microscope. It s about $250,000. But it s not one that you re going to find in every classroom. And certainly you re not going to find in high school environments. But the researcher who works on this actually wants to make it available as a teaching tool. And so what he started by attaching the campus cable TV network, to bring the video into classrooms. And what we did to work with him is to take that same content and digitize it and make it available across the network beyond the campus. So independent of distance, people can sit in front of it and get the high-resolution imagery. This is why bandwidth is important. But also to interact with it with the same kind of feel as if you re actually sitting in front of the scanning electronic microscope. That is, the response time has to be the same as if you were sitting in front of it. So again, the interactivity is important. As you move the stage around as you want the response time to feel like you re actually interacting with it, sitting in front of it. Otherwise, it s not a useful environment. And so this kind of thing he s now taking into high schools that have high speed connections and using it as a teaching tool to bring this kind of exposure to, again, students who might not otherwise have access to these kinds of instruments. Now another example in the remote instrument category comes from the University of North Carolina Chapel Hill, which has a similar model where from hundreds or thousands of miles away you can interact with images coming from this particular scanned-probe microscope. But what they ve added to this environment is force- field feedback. That is, the person is holding a wand in his left hand that he brushes across the tobacco mosaic virus here in this image. Once can actually *feel* the surface of that virus by the information coming across from the microscope. So we ve now extended this from just how I can visualize this at a distance, but now I m adding touch to the application. And when you think about interacting with this -- I actually had a chance to move this wand across some blood platelets and feel it myself. And that s just mind-boggling to somebody who was a Biology Major as an undergraduate, who only understood this stuff either by looking at pictures in a book or maybe looking through a microscope. But actually being able to touch and feel these things down at nanometer scale changes your interaction with the world.

Here we have also some health care examples. This particular one is an example of an individual in a MRI machine undergoing a brain scan. And, in real time, the imagery is presented to the assessing physician to see the brain activity while the person is viewing a visual stimulus. One way to get appreciation for this is to actually hear the developer talk about this.

[BEGIN VIDEO]

The visualization shows the data acquired during our first real time functional MRI experiment. A striking patterned visual stimulus was presented to our subject and dispersed with the presentation of a cross-hair control stimulus. This overhead translucent view of her brain reveals the activity in a region at the back of her brain that is active when she views the pattern stimulus, but not when she views the cross hair. Here we see much activity when the pattern is on. Here we see little activity in the region when the pattern is off. The ability to identify and visualize these regions in real time required 128 processors of the Cray T3E. Real time capability opens up new scientific experimental possibilities and will eventually have dramatic clinical applications, including disease diagnosis and surgical planning.

[END VIDEO]

I look at this and first I note that this application is only possible because of high-performance networks. Not every MRI machine can have the power of the largest Cray supercomputer and the visualization necessary to make this happen. At the same time, I look at this I start asking questions like, What does consciousness mean? I can actually see the activity in my brain. So it s a powerful tool, but it also opens up questions in terms of how these things are used and how we can interact with this.

Now as you heard, teleimmersion is an important aspect of the future of the environments that we re talking about. And those that are familiar with it, again, the example that was given of the *Star Trek Holodeck*, is a useful way of trying to characterize this. This environment, in terms of the CAVE (Cave Automatic Virtual Environment) is a ten foot by ten foot room with rear projection onto the walls and the floors to give you a three dimensional view of the space that you re interacting with. And you ll have a chance to learn more about this. There s a CAVE in this building across the lobby where there will be some demonstrations. The individual here in the CAVE is seeing everything in 3-D. And the important thing about this kind of environment is that you re no longer on the outside looking at a screen. You re actually inside the display environment with the information presented around you. And that information can be things blown up larger than life. Again, it you are analyzing the structures of a cell, you can see the folds of the Golgi Complex apparatus and walk around them and get a feel for those structures (e.g. http://cellbio.utmb.edu/cellbio/golgi.htm). Or as a cosmologist, you can bring the galaxies down to you inside this space and interact with the galaxies, in a space and facility like this, thus truly altering the potentials for interaction with your data.

Going back to the space physics example I showed earlier and the campaign that they had earlier this year, they displayed the data in a CAVE which allowed them to actually see new things in their data that they hadn t seen before as presented as two-dimensional imagery. It s actually opened up insights into new areas of exploration based on these new tools available to the scientist.

Now, a CAVE costs about a million dollars to construct when you add the construction, the computing equipment, the displays, and so forth. More affordable are displays like the ImmersaDesk which is a drafting size version, which costs about $200,000 to utilize. It doesn t give you the full perspective, but it does give you a chance to interact.

Now, I show here on the left-hand side something developed at the University of Illinois, Chicago, a tool to interact with elementary school students, to show them the life cycle of a farm over four seasons. So urban-based students who didn t have a chance to actually go to a farm can now interact with the simulations and with other people in this space. By networking these CAVES and ImmersaDesk together you can interact with other people as if you were in the same room.

Just to throw in another example here, this one comes from Boston University and is called Art World. Here they are actually creating a space for artists to leave animations and clips and sculptures for others to come by and interact with and explore. It s a very compelling and fun application. Of course, none of this is taking away the experience of physical interaction, but it s augmenting and extending the reach of these collaborations.

To give you a flavor for this, there s a clip here that shows interaction among scientists from Virginia with two different locations in Illinois. So, again, spread over many hundreds of miles. People using different interfaces here: CAVES, power walls, which are wall-sized type displays, ImmersaDesks, as you see here, plus standard displays. So to participate in this, again, you need a high performance network connection. But you can also do this from a desktop display.

[BEGIN VIDEO]

Alliance researchers have created a virtual reality application called Cave5D, which enables real time data exploration using the CAVE, ImmersaDesk or power wall. One of the components of Cave5D is Virtual Director, a software framework that allows a user to interactively construct, record, and preview scientific visualizations using voice and gesture input. The other component a configurable application that lets scientists visualize and interact with multi-dimensional numerical data from atmospheric, oceanographic, and other similar domains.

[HURRICANE DATA]

And this ability to collaboratively view and annotate virtual environments gives us the ability to archive these virtual experiences for later use.

[END VIDEO]

So you can actually see people waving. And the faces you see here are actually the perspectives of others in these virtual spaces. And, in fact, since they ve made this video they ve started putting texture mapping of individual faces on here, so you actually can recognize the people.

This application is the kind of environment that we re going to be using for collaborating from our offices, from our laboratories, from our classrooms. And the benefits for this are many. As I say, new ways for interacting with data, with collaborating with each other. Also, one of the individuals in this particular illustration is in a wheelchair. And so while his ability to travel may be more limited than some of the others, his ability to participate in this collaborative environment is not limited. As so, again, using these tools to extend access.

One of the things that we re trying to do with the Internet2 project -- and this is specifically the role that I play-- is to try to bring awareness of the technological capabilities and examples from other areas of inquiry to the faculty through visiting various universities. Last week I was at North Dakota State University. This week I ll be at Virginia Tech, talking to faculty, letting them know what other institutions are doing, and hearing what s going on on their campuses. So they see things, like the space physics example, not as simply a model for a collaborative environment for space physicists, but as framework that can be utilized with their own content and their own interactions as well, and as something that can be brought into their own disciplines.

I ve talked about applications, and I want to just reiterate a point that John Evans made about some of the policy implications. And there s a whole middleware space, that is, the enabling services for these kinds of applications. And I refer to this as a services-rich network environment. And that our focus is moving from circuits to services. Where investment has been in building the physical links among our institutions, we re now trying to provide richer services for the applications developers. And these functional services are going to build the ideal Internet2 where you can build new applications quickly, where the faculty member doesn t have to worry about things like security and directory services and other enablers because the building blocks are made available for him or her. And it s scaled up to the size of the research and education community. The kinds of services we talk about here are things like security including authentication and authorization: who you are and what you can get access to. Things like quality of the service in the network. Who gets priority of their traffic? Does the president of the university get high priority no matter what he or she does? Or is it only if you have video traffic do you get higher priority over e-mail, which can go more in the background. These kinds of services are being built into the network. But how do we actually manage and deploy them in a way that supports our original goals of building advanced applications in support of research, teaching, learning, and clinical environments?

This year I ve been putting a special emphasis on the arts applications. At our October member meeting four weeks from today, we will be doing demonstrations at the University of Washington campus, actually a live performance event, bringing dance and music into a theater setting. The particular installations represented here out in the lobby will give you a chance to interact with some of these people who are building these applications from the creation of music, based on interaction with brain waves; a virtual museum application, among others. Here artists explore how they can still create art for others and to convey the emotions that they may want to share. And they can also appreciate art installations created remotely.

You ll also see some digital library archive materials provided by the National Gallery for the Spoken Word, which are audio recordings of the last 100 years, as well as a video encyclopedia of the 20th century, which are video recordings. And being able to bring these things in high integrity across the network to your desktop no matter where you are.

I will close with a slide that points to more information. Please contact me to find out more about the Internet2 project. I would be happy to supply the URL to get your own copy of these clips and the presentation materials [http://apps.internet2.edu/talks]. And so again, thank you for the opportunity to talk to you, and I look forward to interacting with you. Thanks very much