

From: Teresa Beamer <BEAMER@DENISON.EDU>  
Subject: Minutes from April 11, 2001 OARtech meeting  
X-cc: dunet@denison.edu  
To: OARTECH@LISTS.OAR.NET  
Status: RO

OARtech 4/11/2001

## Introductions

Akamai - Patrick Gilmore, Kirsten Fitzgerald  
Antioch University - Bruce Friend  
Case Western Reserve University - Eric Chan  
Clark State Community College - Hugh Evans, Jim Gossett  
Denison University - Teresa Beamer  
DeVRY Institute of Technology - Dave Leitch  
Edison Community College - Chuck Jones  
Heidelberg College - Kurt Huenemann  
Kent State University - Ransel Yoho  
Lakeland Community College - Dave Levine  
Lorain County Community College - Lisa Guerrero, Norman Lease  
Medical College of Ohio - Chris Bauer, Joe Kitting  
Miami University - Tim Gruenhagen  
NEOUCOM - Bill Mayhew  
OARnet - Christopher Cook, Ruth Crites, Fred Crowner, Linda Roos, Albert  
School, Pankaj Shah, Kelly Sitz, Gene Wallis  
Oberlin College - Cal Frye  
Ohio State University - Mowgli Assor  
OhioLink - Anita Cook  
Shawnee State University - Mike Pinson  
Sinclair Community College - Roz Taylor  
University of Akron - Debra Keller  
University of Cincinnati - Kim Koeppe  
University of Dayton - Frank Tavenner, Ron Wagers  
University of Rio Grande - Kingsley Meyer, Mike Snider  
Wright State University - Larry Fox, Patricia Vendt

Old Business - The wireless demo has been rescheduled for the June meeting. Do we want multiple vendors in on one day? The discussion said that they are interested in looking at what differentiates each vendor. Have an 802.11 overview, then find out what the problems with it. Then split the time between vendors to see how they are different.

## Doug Gale, - Dark Fiber Infrastructure Project

Doug reviewed briefly why OARnet is interested in the project. OARnet is doing things to buy us time until we can manage bandwidth. Use keeps going up faster than costs are dropping. Local loop costs are not going down. OARnet is addressing the intrastate issues and local loop costs by participating in the Mission fiber project - 2 strands of fiber throughout the states with 33 pops. Schools could participate by purchasing fiber to connect to the closest pop at cost +10%. OARnet would be running each pair as multiple OC-192. A spreadsheet of how the project would be financed was handed out. OARnet will own the fiber and would be able to use it as they needed. They could use a wavelength for supercomputing, another for video conferencing, another for medical uses, etc....

The map shows Michigan and Pennsylvania. Would they be participating in the project? Mission is looking at those states but is implementing in Ohio first.

What is the risk? When the contract is signed OARnet must put down 20%. If the project dies completely, it would wipe out the contingency fund. Would it effect our rates? Not expected to. Another risk is that if OARnet doesn't do anything and things continue on the current path, your costs will continue to grow.

Butler County was looking at adding a fiber network within the county. OARnet recommended that they contact Mission. Mission will be pulling the fiber and will be able to put a loop in the county and Miami will be able to connect to it.

What is the Timeframe? The signing must occur this month or it will probably not occur. The company is getting impatient with us. The process is currently working through the Ohio State political channels. Roll out is 18 months. Once they have a signed contract OARnet will let us know what the implementation schedule is.

Do you feel we'll make the end of the month deadline? Doug feels we have a 70% chance.

## Akamai Presentation

Slides are available at <http://www.oartech.oar.net/library/presentations/>

Kristen Fitzgerald, Patrick Gilmore

Akamai is based in Cambridge, MA across from MIT. Founded by a MIT

research team including F. Thomson Leighton, Chief Scientist, and Danny Lewin, (current CTO). It was seeded in 1995 as research effort to improve the Internet content distribution problem. Currently they have 3800 customers.

Akamai is a Content Distribution Network (CDN). A CDN has 3 main components - edge caches, content delivery, and mapping. Akamai is not the only CDN, but is currently the largest. They have 9700 servers all over the world. They work with ISPs and networks all over the world. The other issue is getting content to the edge including object, videos, etc. Mapping is Akamai's forte. Finding the closest edge server for each user. Network proximity does not necessarily mean geographical proximity.

Akamai's CDN is hosted on an "origin server" that is accessible to all edge caches. They do not provide the origin host site. They are currently working on a hierarchy of the caches. Users only speak to edge caches, not the origin server, to get distributed content. The edge caches pull content from the origin server then delivers it to the many end users. This is a nearly infinite scalability model. It avoids congestion and long latency and provides reliability and redundancy. Different servers do different things, some do video streaming and some do data. If a server is determined down, their system redirects traffic after 20 seconds. They generally give their servers to ISPs, educational institutions, government, etc....

Their smallest install is 3 servers. Sites that install don't need to change their network. It's just another host on your LAN. They have 24x7 NOC and can help you with other network issues as well. They pay for the cost of shipping etc....

Region - a set of edge caches. A smallest region would have 3 servers. The measure of large versus small is based on the amount of traffic they would be looking at. They have found that students tend to do twice as much traffic as the typical dialup users because they tend to run 24x7 at 10/100 base T. For a school of 10,000 students they would recommend a five-server configuration. Servers are Linux based using shrink-wrapped software. Your region serves only your users.

If you have pipes that go to other places, then they will pay to have you provide the cache for those other places. The servers are dual Pentium IIIs with a Gigabyte of RAM with 2 or 4 drives and 2 100baseT-network cards. Also they ship an Ethernet switch with 100baseT or gigabit uplinks as well as the patch cables. Configuration diagrams are emailed out to you. Sites will only have to bring it up and include the IP and Akamai will take over.

Each server has 2 public IP addresses, a service address (virtual) and the physical address. If these addresses are down, they consider the server as down. The servers come in pairs as buddies. If a server stops responding, and the buddy determines that it's buddy is down, then it will begin answering the current requests. Once their mapping system determines that a system is down, they redirect all traffic away from both servers.

## Content Delivery

Possible bottlenecks occur at the first mile, backbone, peering points, and last mile. When the user requests the content, they are directed to the nearest cache. If it cannot find it on the current cache the cache asks its buddy server. If that doesn't have it they go out to the Internet to get a copy.

How frequently you download the information to determine the currency? They check for a fingerprint using meta-data to verify that they have the most current page. The least used pages are expired off the servers. The Akamai servers have been spoofed to get non-Akamai content.

How do you interact with existing caches? They ask that you not cache the Akamai servers on your network.

Do the Akamai servers need to be inside or outside the site firewall? They want their servers outside the firewalls. They don't want you to trust their servers. They also use odd ports for maintenance and firewalls tend to stop that.

The customers include the caches in the html code that tells the user to download the gifs, etc... from the Akamai servers to download just the objects. This product is FreeFlow - it serves objects. Note that axxx.g.akamai.net returns different addresses depending on where on the Internet the lookup occurs.

EdgeSuite caches the whole site. Cnames redirect the lookup to customer.d4p.net, then sends customer.d4p.net to axxx.g.akamai.net. This means the end user never talks to the origin server. They always talk to the cache. Uncacheable content is tunneled back to the origin.

Video on Demand uses http to get to the content to the edge caches. Content is streamed in native format from the cache to the user. Cache only downloads from the origin the frames needed by the user.

Live uses UDP to transmit content to the edge caches. They translate all

streaming formats into UDP then translates it back to the native protocol for the end user.

Live "SteadyStream" uses top level reflectors to provide redundancy for the streams. Multiple streams are sent to the region and the cache picks from the best stream. This allows global distribution of real-time video and users can connect anywhere and get good quality video over the public Internet. The problems are that each region must have 3 users to have effective use of the bandwidth. They do QuickTime Streaming downloaded via HTTP.

They maintain core Hierarchy Regions, super caches. Each edge cache can go back to the super cache to get info. The core region can go back to the origin. The origin only has to deal with the core region. This avoids the TCP slow start connection because they have persistent TCP connections. The origin server maintains the TCP connection with their cache.

Mapping algorithms direct the end users to connect to the nearest edge cache using DNS. The main components in the algorithm for finding the closest edge cache are packet loss, throughput, and latency. They send the user to the server that contains the content. They use a special hashing algorithm to find it. They use sub-delegation to return the IP for a DNS lookup to find the IP for the Akamai server in their area of the net. Mapping is based on the local name server's IP address, not the end user's IP address. They map (refresh) the network conditions every 15-20 minutes. Addresses returned are given a TTL of 20 minutes. The low level name servers are on the edge caches. They poll the others within the region every 3-20 seconds for availability, load, etc.... They always return 2 IPs. The final IP Address TTL is 20 seconds.

Akamai – a-kuh-my (Hawaiian) means clever and slang for cool.

What does Akamai expect in return for installing their servers? Power, T1.

What do the schools benefit by participating if OARnet is running them? Akamai would like their servers to be ubiquitous. This puts the cache closer to the users. They ask that you provide access if they need it. They also want to work with you on the filters.

Why would schools not do it? Rack space limitation.

Statistics - there is a web page that customers can look at for information.

Time for installation: Once the contract and technical questionnaire is turned

in. Then the order for the equipment is placed. The order takes 4-6 weeks. The equipment comes in. The school installs it and sends back an installation record to them. Their installation group does the global install and brings them into the map. The whole process takes 4-6 weeks.

During the installation process if the server does down, then server install will be delayed. But users will be redirected to another server.

Can the site disconnect the server if they suspect the cache is causing a problem? Call the support line, they will redirect the traffic if necessary. In the end, this is your network and you can disconnect the servers, but please let them know what is happening.

Ohio sites that currently have Akamai - OSU, OARnet, Denison, and Ohio Wesleyan. What schools are customers of Akamai? MIT, Johns Hopkins, U of Notre Dame, etc....

Lunch

Whitepaper.

When presented to them, Osteer said they would tentatively approve it but requested that we provide a description on the bandwidth requirements for protocols, especially in the area of distance learning delivery media. We are looking for a copy of the bandwidth requirements of certain protocols that was distributed when the whitepaper was originally crafted. Osteer would also like a section on PDA recommendations. There was some discussion on how this could be done since the PDA market is all over the map as far as hardware and OS. OhioLink has an application that does use the palm. Anita will email the requirements to Bill for that application as a starting point. Bill will draft a section and send it out to the OARtech list for discussion.

Osteer wants a bandwidth requirement graph like the "Cisco" graph/per user bandwidth graph. They would like this as a justification section for bandwidth needed. It was agreed that it should include other apps besides multimedia applications. Please email possible applications and the bandwidth needed to support those applications to Debbie Keller, keller@uakron.edu. She will try to summarize the information into a graphic. We will need to make the labels understandable by a manager. These are per user at the LAN level. Bill found [www.mpeg.org](http://www.mpeg.org) has some comparable standards information that we may be able to use. Multicast issues may have to be addressed as a footnote. Be watching the OARtech list for the discussions and please participate.

Academic Services Update

## Mogli- OSU - Security

Mogli summarized several security issues that OSU has been seeing. These include an NTP bug, proliferation of worms aimed at Linux boxes, exploitation of Bind (name server), and an LPD worm.

Someone got VeriSign to give out a class 3 certificates that are labeled Microsoft Corporation to someone not associated with Microsoft. A patch from Microsoft is available.

New staff.

Christopher Cook - new services and marketing

Lindsay Roos - Internet 2 person

Group sponsored I2 application from OARnet for OARnet member schools was sent out. I2 has sent back a request for more information. They wanted more information on what the uses of I2 will be implemented within the group. There is currently no timeframe known for the application. How would the connections be made? Would it be able to use their current connection? Sponsored educational groups are called SEGPs. We would be interested in seeing what is currently be used on I2 for current Ohio members.

Christopher Cook passed out a listing all the services currently offered. Chris has a survey for the directory services workshop so they can better determine the content of that class. If you go to [www.oar.net](http://www.oar.net) and go to the academic link, you'll find a list of the online courses available. You and staff can look at it and review to see if it would be useful for your group. Mind Leaders offer the courses. The costs are posted on the web site.

Window on the Future - sponsored by ITEC

Kelly Sitz

Program coordinator for ITEC-OHIO

They had over 120 attendees from academic, government, corporate partners, and corporate labs. It was a forum for individuals to discuss currently emerging technologies. Thursday night they had a keynote. Friday was various faculty presentations. Looking to add more corporate partners (current 8). They are in the process creating an ITEC steering committee, [Sitz@oar.net](mailto:Sitz@oar.net) for anyone interested in the makeup of the steering committee. Next Conference will be in December.

Current partners include Spirent, Polycom/Accord, Web Sprocket, etc... generally look at emerging technology partners. There is a membership fee for

ITEC but it is waived for the Ohio I2 members - so once the application is approved OARnet schools will automatically be members.

Albert School - Security group update

Looking at what technologies we need to be looking at 1-2 years out. They are also selling some services.

Currently working on what PKI is happening to interact with the government as the PKI infrastructure immerges. Currently working with I2 - 2 possible certificate authentication models - CREN and I2. It looks like there will be a bridge certificate authority to provide the communication between the two types. By 2003-4 the main government entities will not accept paper documents anymore. So these entities will need electronic signatures. The federal PKI info has put up their bridge certificate authority and is operating. They have done a years worth of testing and working. The bridge certificate authority provides inter-operating between different types of PKI. Csrc.nist.gov/pki has the information on what is happening on PKI.

OARnet is in phase 1 of a PKI pilot. They are primarily coming up with boilerplate PKI certificates and interoperability testing. There is some Microsoft software coming out that does certificate authority functions that may not work with Windows 2000. Based on x509 certificates, using client side certificates. Would be interested in keeping in touch with what it happening on each campus. The first agencies in the government that will have to have this in place are the universities. The information on this requirement would come through your research and initiatives representative. OARnet is trying to develop something that will leverage the members needs and keep it within the cost level that would be doable by the individual sites. Albert will be sending updates to the OARtech list. His address is aschool@oar.net. To send to the OARtech list send to oartech@lists.oar.net.

Gene Wallis

Question on how much reduction in Internet bandwidth is seen with Akamai caches? OSU sees about 1% reduction in bandwidth use. Denison sees anywhere between 4-10%.

OARnet currently has 20 servers. 13 are http servers, 7 are streaming servers. OARnet is also selling bandwidth to Akamai. Has OARnet had any troubleshooting issues? You don't know where the traffic is coming from. So even shutting off the server didn't clarify the situation because it came from other servers. They have not been able to "prove" any problems caused by them. They have had to replace about 1 a month due to age. Current servers seem more reliable.



What are the cons?

DeVry was having a problem getting to Yahoo.com a prime customer for Akamai. What is happening? Akamai was innocent on this problem. Gene indicated that Cisco's IOS has a problem that can prevent the routers from passing large packets - 12.0.13.6 IOS and above. OARnet now knows the symptoms and knows how to correct it when they see it, but Cisco is working on the problem. It's a TCP problem and sets the MTU to some number. The most common is in the 700-bit range. The problem is address dependent. This problem has bit several of the other networks as well as OARnet.

Would like something to let us know when something like this is happening. It has had major ramifications as far as local troubleshooting issues. This started occurring the last week in March.

Another problem that has occur recently: OARnet brought up a new router in Cleveland and every 6 hours the router would go down, replaced hardware, IOS and now have an answer - the new hardware version of the pa-a3 (ds3 interface for that router) is different from the previous one. When you go to that router and do a show version, it causes a chain reaction that causes the interface to go down. OARnet has software that goes out every six hours to check the router and was causing the router to go down. They are replacing the card with an older one that doesn't have the problem.

DeVry had a problem where MRTG shows a 600-mb flat line on Internet link. With discussion, it was suggested that the snmp daemon might have quit responding.

Another problem that happened last night: A site called in couldn't get to their web server from outside their network. The site had upgraded their PIX firewall. It was answering for everything on the network.

Who should we contact to discuss how they can connect to I2. Technical issues should be addressed to Gene.

AEP project - proceeding slowly. They have brought up 2 pops. Portsmouth is live at this time. Other sites have equipment waiting for the lines to be installed.

Approval of the minutes - passed.

Tim Gruenhagen, Miami University was nominated for vice Chair. Ransel has volunteered to continue as web master for the web site. Any information that would be appropriate for the web page (e.g. past presentations, minutes,

etc...) can be sent to oartechwm@oar.net, or directly to Ransel  
ransel@kent.edu.

Meeting adjourned.