ITU and IETF in Agreement on ENUM Administration

Letter from ITU To ICANN Blocks .tel gTLD Applications As Competition to ENUM -- Administration Modeled on Neutral Tier 1 Database Holder of Pointers to Records of Provisioned Services

Editor's Note: The ENUM process has taken some significant further steps towards implementation since the publication (in mid October) of our December issue that contained the ENUM interview with Richard Shockey of Neustar. In this brief article we highlight the most recent developments. The material that follows is based on conversations with Shockey.

According to Shockey, "communications between the IETF as a professional engineering organization with the ITU Study groups which are trying to solve engineering problems is in fact excellent and on going on any number of fronts. However problems tend to come when you have to deal with the ITU Secretariat which is the organizations political side and has its own view of the world."

"What happened in Berlin with the SG2 meeting with the IETF that concluded on October 26 was in my judgement a raging success. What happened was a recognition by the ITU that its E164 standard is a good thing and that ENUM itself can be modeled along the lines of E164 in a successful manner that does not compromise the security and stability of the Public Switched Telephone Network. And that furthermore that ENUM also can be deployed in a way that deliberately respects the rights and prerogatives of nation states involving telephone numbering."

IETF and ITU in Agreement on ENUM

On November 1 a document entitled "Liaison to IETF/ISOC on ENUM" and authored by the ITU-T Working Party 1/2, in Berlin, between 19 - 26 October 2000 was published on the IETF Announce list and ENUM working group list. The document conveys the understanding by ITU Study Group 2 of how it agreed to implement E164.arpa in meetings with the IETF in Berlin between October 19th and 26th. Shockey added: "As of November 10th there had been no complaints on either the ENUM or the IETF discuss list. This is usually taken to mean consensus. Therefore it looks as though at this point the path is clear to proceed towards implementation."

For the Liason document itself readers should turn to:

Title	: Liaison to IETF/ISOC on ENUM
Author(s)	: R. Blane
Filename	: draft-itu-sg2-liason-enum-01.txt
Pages	: 7
Date	: 08-Nov-00

Abstract: Working Party 1/2, of the International Telecommunication Union P Telecommunication Standardization Sector (ITU-T) held a meeting of its collaborators in Berlin Germany 19-26 October 2000. The agenda of the meeting contained several contributions regarding RFC 2916: 'E.164 Number and DNS' from the Internet Engineering Task Force's (IETF) ENUM Working Group - more specifically, the method for administering and maintaining the E.164-based resources in the Domain Name System (DNS) as related to the ENUM protocol. Consequently, in addition to the WP1/2 collaborators, there were several members of the IETF present to assist with the discussion of issues contained in the aforementioned contributions

A URL for this Internet-Draft is: http://www.ietf.org/internet-drafts/draft-itu-sg2-liason-enum-01.txt

Shockey continued: "It is important that your readers understand that these developments have been driven by the IETF, ITU and DoC independent of ICANN." From various conversations with knowledgeable sources we have been assured that ICANN has been told what to do and so far has complied. First with the written instructions form John Klensin on behalf of the IAB to Mike Roberts to insert e164.arpa into the ICANN root in September. The second step occurred in October when four different organizations put forward gTLD proposals to ICANN for .tel or variations on the same.

The implementation of the .tels would essentially duplicate the functionality of ENUM. Such implementation would leave registrants for ENUM services with domains that could, according to ICANN's own rules, be hi jacked since ICANN owned the name and not the customer. This move clashed head on with the path taken by the IETF and ITU on ENUM. The ENUM folk did their lobbying.

On Novemeber 1, 2000 in a letter <http://www.icann.org/tlds/correspondence/itu-response-01nov00.htm> to Mike Roberts, Yoshio Utsumi, the ITU Secretary General told Roberts to back off and not to inaugurate any names that would compete in any way with e164.arpa. After various paragraphs of diplomatic circumlocution the letter concluded: "As I am sure you are aware, the E.164 international public telecommunication numbering plan is a politically significant numbering resource with direct implications of national sovereignty. It is subject to a multitude of national approaches, regulatory provisions, and, in some cases, multilateral treaty provisions. Considering this, governments should be given the opportunity to fully reflect upon how their particular numbering resource responsibilities relate to DNS-based telephony resources."

"In this regard, the ITU is working with the IETF to progress a careful exploration of these complicated issues in the context of its joint work concerning the ENUM protocol. As there are still considerable areas of coordination work needed at this time, until there is an opportunity to further explore the issues within the context of joint work underway and with national governments; it is the view of ITU that it would be premature for ICANN to grant any E.164-related TLD application as this may jeopardize these cooperative activities or prejudice future DNS IP Telephony addressing requirements."

In other words a strong warning not to do anything to harm the usability of the e164 numbering scheme that had been picked for ENUM. The letter was in early another warning shot across ICANN's bow that that the IETF and ITU not ICANN would write the rules for an ENUM implementing TLD. The result is that with the full support of the IETF and ITU ICANN has formally been told by the secretary general of the ITU in strong diplomatic language to back off and not approve any of the four dot tel poposals. Therefore, in view of the IETF's declaration in the spring of 2000 that .arpa was a gTLD to be used by the IETF for infrastructure purposes, it seems reasonable to look at E164.arpa as ordered into the ICANN root by John Klensin and Karen Rose in September as a new TLD that is not under ICANN's control.

On November 10 ICANN staff released its evaluation of the new TLDs. What the document had to say about the .tel applications showed that ICANN had gotten the ITU message. "Based on application of the August 15 Criteria, the evaluation team believes that none of the four proposals in the telephony-related group should be selected at this time. Each of the four proposals appears not to have adequately addressed requirements for stable, authoritative coordination with the PSTN numbering system, particularly when dynamic-routing considerations are taken into account. (Of the four, Group One, Number.tel, and Pulver/Peek/Marschel are of particular concern in this area.)

In addition, the Group One Registry, Number.tel and Pulver/Peek/Marschel proposals would do little to address unmet needs. Moreover, if a TLD were established in which the service available at URLs was defined by the TLD rather than the prefix, this would likely increase confusion regarding URL naming conventions. Finally, the concerns raised and caution urged by the ITU counsel against establishing a telephony-related TLD until further study and consensus-building within the Internet and telephony technical communities." Editor: the preceding is accessible at http://www.icann.org/tlds/report/report-iiib3-09nov00.htm

ENUM Administration in North America

What follows is our of the current thinking on the proposed administrative structure for ENUM administration. (These paragraphs are based on notes from a conversation with Richard Shockey.)

In each nation the national telecommunications regulators will be called upon to choose the entity that will actually hold the ENUM resource records for a telephone number. The regulators will also have to determine the rights and responsibilities exercised by that holding entity. The public will have to deal with the entity but will have the ability to ask the regulator to step in and provide due process protection for their rights should they be abused by the entity to which the regulators have granted administrative authority.

In the United States it looks as though there will be a two level system. It is expected that there will be a single Tier 1 entity that will delegate responsibility records and will keep a data base identifying who is authorized to do the provisioning of ENUM types of services assigned to every phone number. The single national tier one entity is there to certify that person x has two things. The first is authority over phone number y. The second is the entity that this person has selected to do the actual service provisioning.

Now tier two ENUM services are to be provided by the owner of the phone number. If that owner is a large entity like a corporation or a university, it may designate its own telecommunications people to provide the services. However not everyone will really want to do this nor will everyone be technically capable of doing it. Therefore it is assumed that companies will spring up to provide these services for a fee to individuals, organizations and small business. The administrative system will be set so that the Tier 1 entity will be the sole trusted provider of data base pointers to the legitimate owners of numbers and the services they have authorized for them.

From a social and economic point of view telephone service is an absolutely critical service for each individual. Indeed ones phone number is becoming a critical global identifier and the government must guarantee that each individual with a phone number will have control both over it and the services provisioned for it. Or in the event the phone number must be changed the guarantee must be for the seamless change of provisioned ENUM services to the new number.

Because ENUM is becoming service control point for Internet telephony in the same way there are service control points in the PSTN, it is intended that in each nation state there will be a single entity identifying the authorized provisioner of the resource records. These resource records are the means of ultimate control. The tier one entity is to be the repository of the services subscribed to by the owner and user of each phone number. The repository has authority to change its records only on receipt of validated instructions from the number's owner or the tier two agent that the owner has authorized to act on its behalf. With ENUM this puts control of ENUM and internet telephony services into the hands of the subscriber and allows intelligence and service logic to reside in edge devices of the internet as the quintessential STUPID Network.

ENUM administration will put control of service provisioning for an individual's phone number in the hands of entities responsible to the interests of that individual though a process of government ensured accountability and due process. The COOK Report concludes that the proposed process is deliberately unlike that of ICANN which assumes ownership of and total control over a domain name. The process developed has been designed to keep control over ENUM out of ICANN's hands. In other words the administrative model just described has been structured to remove it form the ICANN control and place it into the hands of a regulator responsible to democratically controlled government who will give the owner phone number full control of services attached to such a number. Policy is to be set by the regulator as a part of the political process within each national numbering system. Everything else is administrative. New policy involving the ownership of a number and the control of a resource record cannot be set without the regulator's approval. Other policies may be set by the tier 2 provisioning entities just so long as national policy on ownership and control is not overturned.

In the United States it is expected that the locus of policy will be either the Department of Commerce or the Federal Communications Commission. An early decision of the new administration will be to authorize either agency to issue an RFC for private sector companies to offer bids on how they would design and implement a single Tier One entity for the United States. Responding to such a RFC will take most of 2001. In the meantime setting up and operating test beds for ENUM services is a major priority.

For a similar point of view (minus the ICANN references) readers should turn to

Title	: ENUM Administrative Process in the U.S.A.
Author(s)	: P. Pfautz, J. Yu
Filename	: draft-pfautz-yu-enum-adm-00.txt
Pages	:
Date	: 18-Oct-00

Abstract: This document considers administrative processes for ENUM in the U.S.A. and offers two 'strawman' proposals in the spirit of moving forward the work that must be done to implement a useful ENUM capability. The U.S.A. has implemented number portability; therefore, it is the telephony user that controls the assigned telephone number so long as it maintains the telephony service. While the proposed processes are tailored for the U.S.A. they may be appropriate for use by other countries that implement number portability so that the donor telephony service provider (e.g., the telephony service provider that is assigned a block of telephony numbers before any number porting event happens from that number block) is not relied on for maintaining the delegation information for a telephone number (e.g., the Tier 1 function in the ENUM process).

A URL for this Internet-Draft is: http://www.ietf.org/internet-drafts/draft-pfautz-yu-enum-adm-00.txt

COOK Report January 2001

Klensin Internet Drafts Propose Radical DNS Revamp - New Class Means New Root

Drafts Are Aftermath of Network Solutions Split With IETF on ENUM and Internationalized Domain Names

Part One:

ENUM Trials

Changes are afoot in DNS and ENUM. VeriSign-Network Solutions announced the opening of the 'first public ENUM trial" on December 20th. The talk was of the development of applications and hyped the fact that the trial would rest on the "DNS infrastructure [.com, .net and .org] owned and managed by VeriSign." It implied that it was VeriSign ownership of the .com database and expertise in security that would make ENUM work in a way presumably superior to anyone else. In making these assertions Versign - Network Solutions again started writing its own version of reality refusing to admit that (1) we are dealing with an open IETF standard protocol subject to implementation by anyone, that (2) everyone has equal access to the .com infrastructure and that (3)VeriSign type security is not uniquely meaningful to the implementation of the protocol. In order to insure security in ENUM, one must implement DNSSEC. DNSSEC does not rely upon VeriSign implemented X.509 certification. It can also be argued that it isn't really an ENUM trial, as ENUM, by definition, resides in e164.arpa..

Instead of trying to make a global standard Net Sol is determined to establish a competing proprietary approach to meet head on with a globally agreed IETF-ITU standard. This is an action that is contrary to everything that has pushed the Internet forward to its present position. The motivation for this may well be what an observer called the biggest land grab in history of the DNS — the goal being one for Network Solutions to jump out ahead of the market and by dint of shear numbers garner enough support to have a de-facto standard. Certainly it can be argued that at a minimum that Network Solutions-VeriSign's actions by staving off the IETF - ITU approach prevents it from becoming a fait accompli and leaves other options open. Such options however do little for the Internet using community at large. For, it can be argued that if ENUM is going to be implemented, e164.arpa is the stem to use. Using e164. with com,net,and org gets you into the same situation as dealing with multiple root name servers (a risk of inconsistent mappings), and you're no longer implementing ENUM, but rather, an ENUM-like protocol that may conflict with ENUM.

The NTIA ENUM Rountable

Prior to December 18, 2000 <http://www.ntia.doc.gov/ntiahome/ntiageneral/enum120800.htm > NTIA announced that on 12/18/2000: The National Telecommunications and Information Administration (NTIA), U.S. Department of Commerce, would hold a roundtable to discuss and explore issues related to the Internet Engineering Task Force's (IETF) Telephone Number Mapping (ENUM) protocol and the work being undertaken between the IETF and the International Telecommunication Union (ITU) Study Group 2 (SG2) to consider how number resolution using ENUM may be affected by public switched telephone network infrastructure and telephone numbering administration.

The official presentations from the meeting are found at p://www.ntia.doc.gov/oiahome/enum/doc_plp.ppt" http://www.ntia.doc.gov/oiahome/enum/doc_plp.ppt They included on the IETF/ITU side Richard Shockey and two presenters from ATT - one of whom was Penn Pfautz author of an internet draft on ENUM administration that we have discussed before. Tony Rutkowski presented for VeriSign and Network Solutions.

According to some attendees of the meeting VeriSign, Network Solutions and its partner Telcordia were obfuscating the issue of what ENUM is. They ignored the issue of administration and the right of multiple application service providers to offer ENUM services on behalf the owner of the phone number. According to the Network Solutions version of reality explained at the meeting and in the press release cited above ENUM is about putting phone numbers into the .com, .net and.org and perhaps into other gTLDs as approved by ICANN.

Apparently Net Numbers and all the other the dot tel gTLD applicants that got shot down by the letter from the ITU to ICANN's Mike Roberts were at NTIA aligned with VeriSign and Network Solutions on doing a DNS approach. By going through ICANN gTLD licensing procedures, they were ready to deliver the consumer into a "free market" solution where the consumer's rights would be dependent on ICANN's good graces in exchange for package where Network Solutions could in an interesting land grab assure that it would be in on at least ENUM registry fees for the world's 14 million dot com addresses.

Anyone who has read RFC 2916 and the related IETF ITU documents would know that ENUM is not about putting phone numbers into .com, .net and .org. Rather it is about putting phone numbers into e164.arpa. Sources tell us that Rutkowski on behalf of VeriSign NetSol derided e164 dot arpa as creating an unneeded government monopoly and would impose on the US the unwarranted burden of accepting an ITU protocol (e164) when there was really no need to do so. Another observer had a slightly different point of view: "The US already accepts the E.164 protocol. Administration of telephone country code "1" has been delegated by the US government to Neustar. The concern is that somehow (and I'll admit being a bit vague as to exactly how), the US government mandating an administrator for the 1.e164.arpa registry would cede control of (something) to the ITU."

Of course in doing so they ignored that the rest of the world including perhaps Canada would likely go with national administrators which kept databases of which service providers the owners of phone numbers had delegated responsibility for delivering a variety of ENUM services. As a result Americans and other dot com owners can register numbers in .com but if rest of the world goes e164 then they will be invisible. Why? Because if the rest of the world stays e164, you will have two incompatible classes of look ups.

An expert whom we asked to review our draft replied: "Not necessarily. Remember that essentially no applications exist which make use of ENUM (or an ENUM-like protocol). One possible approach would be to create an "ENUM path", that is, a list of suffixes to append to a ENUM-ized telephone number for lookup. For example, it is conceivable that there would be something like:

% setenv ENUMPATH :e164.com:e164.org:e164.arpa:telephones-r-us.de:

The application (SIP gateway, WAP phone, whatever) would then try looking up the number 1-650-779-6003 as

getnaptrbyname(3.0.0.6.9.7.7.0.5.6.1.e164.com)
getnaptrbyname(3.0.0.6.9.7.7.0.5.6.1.e164.org)
getnaptrbyname(3.0.0.6.9.7.7.0.5.6.1.e164.arpa)
getnaptrbyname(3.0.0.6.9.7.7.0.5.6.1.telephones-r-us.de)

The application could stop at the first hit, or it could go through all in the path and present the user with a selection to choose from (if choosing is necessary — the proponents of such an approach argue that the various ENUM-like providers would work out a way to avoid assigning the same name to different entities). Conceptually, this is equivalent to dealing with multiple DNS roots, with all that implies.

It is probably also worth noting that countries outside the US would probably be a bit annoyed if a US company took over control of the mapping of E.164 address space into the DNS for that country."

In any case, Networks Solutions has already convinced its VeriSign owners to sell a balkanized and marginally worth while service. The VeriSign - Network Solutions - Telcordia approach is also one that would further marginalize the worth of ENUM services by putting them under a DNS controlled by ICANN where there only protections exerted on behalf of anyone are actions that serve the interests or large corporate trademark interests.

The Network Solutions VeriSign presentation by Rutkowski was seen as a move designed to show the US government that there was no unanimity in the industry and therefore that neither the FCC no DoC should issue an RFP for ENUM Administrator in 2001. Another observer disagreed: "Pragmatically speaking, there actually is no unanimity — very few people have the slightest clue as to what this is about. As such, it is unlikely FCC or DoC will be in any position to decide anything on this issue in the near term."

Given the chaos that has developed around what was assumed to be an IETF - ITU consensus, we conclude that in this instance, we will see again how industry self regulation will become, as Michael Froomkin pointed out in his long legal study of ICANN, industry self-dealing. Protection is available for the interests of inside power groups. Everyone else be damned.

ENUM in Geneva in January 2001

Meanwhile from Robert Shaw we have the following announcement: On 17 January 2001, ITU will host a one-day ENUM Workshop on administrative issues related to deployment of the ENUM protocol. The meeting will be held at ITU Headquarters in Geneva, Switzerland. The following related materials are now available on the ITU web site:

ITU ENUM Web pages
http://www.itu.int/infocom/enum/

Description and Objective of the Workshop
http://www.itu.int/infocom/enum/workshopjan01/enum-ws-jan17.htm

Draft Agenda
http://www.itu.int/infocom/enum/workshopjan01/enum-agenda.htm

Although the workshop is principally intended for national regulators, policy makers and numbering authorities, attendance is open to any interested participant. As we face the January meeting a few of many interesting questions are: What positions will the Canadians take with regard to their share of "1"? Will American nxx codes be put into e164 or into some new GTld or even into dot com?

On September 27 in email to NSI, and officials of the DoC, the IETF, and the ITU Patrik Faltstrom author of key ENUM drafts and IETF Area Director thoroughly castigated Network solutions for announcing ENUM trials accusing it of setting back the cause of ENUM deployment by several years. As contention and conflicting approaches continue to grow it looks like Faltstrom may well have been right.

Part Two:

But the More Critical Issue is the Multi Lingual Testbed

On August 22, 2000 VeriSign NetSol announced that it would open a testbed for the registration of international domain names in not Latin, non ascii character sets. It announced the use of the RACE encoding scheme. On October 31 it announced that as of November 10, it would begin accepting paid Chinese, Japanese and Korean registrations for com, net and org in its testbed. Effectively jumping the gun on an IETF standards process that had not yet decided how to proceed. Their press release claimed that 'VeriSign GRS's approach is to ensure current and future compliance with evolving internationalized domain name standards developed within the Internet Engineering Task Force.

On November 8, 2000 Network World Fusion <http://www.nwfusion.com/news/2000/1108overseas.html?nf> reported that 'several of the world's largest domain name registrars will start selling multilingual domain names Thursday night, [11/09/00] despite a recommendation from the Internet Society that this initiative be delayed because it will harm the stability of the Internet's Domain Name System." "The Internet Society ``believes the commercial testbed is premature under the technical standards of the Internet," the press release states. It recommends that "before users of the Internet establish claims of particular domain names, it is now time for a hiatus in the commercial deployment of internationalized character sets in the operational DNS."

On November 17 in an article headlined Beijing Moves to Control Domain Names <http://www.techweb.com/wire/story/reuters/REU20001117S0001> CMP Tech Web reported: "The Chinese government has mandated that only a handful of domestic companies may assign Chinese-language Internet addresses, striking a blow to the registration service launched last week by VeriSign.

"The move pits VeriSign Inc. (stock: VRSN) and other foreign Web address companies against government-supported Chinese companies in a fight for a new but lucrative market to register addresses in Chinese script. The two sides have launched com-

peting standards for Chinese-script domain names as thousands of confused companies rush to register with both sides."

"Without prior approval, no organization or individual is allowed to manage, provide services for, or act as an agent for Chinese-language domain name registration," said a notice from the Ministry of Information Industry (MII) seen by Reuters on Friday. The notice, published on the Ministry's website, vests the China Internet Network Information Centre (CNNIC) with sole authority over registration of Chinese-character domain names. " In effect the Chinese may be seen as rebelling against the idea that any western company should be able to charge Chinese business annual "rents" on Chinese words used in Internet commerce.

Our reviewer commented: "This is one view. Another view is that the Chinese government is attempting to assert control over anyone who wishes to use Chinese characters in a domain name, regardless of whether they are subject to Chinese governmental influence. For example, a person in Singapore, wishing to register his/her family name in Chinese characters would (one of Singapore's national languages), if the world were to abide by the MII's pronouncement, have to go to a company authorized by the Chinese government. More interestingly perhaps would be the case of a company in Taiwan."

By the end of November Network Solutions was continuing its registration as were several other companies in Asia.. This while the Chinese government steadfastly asserted that only designated Chinese companies had the right to do so. Then on December 8 *Computerwold*, Hong Kong announced that I-DNS.net (a company formed by individuals that participated in work promoted by APNG investigating the feasibility of non-ASCII domain names) registered 90,000 Chinese names in its first four days in business using a system and database completely independent of VeriSign - Network Solutions.

What has happened in the last four months assures that ICANN will face conflicts in getting multi lingual domain names into its legacy root. Should it choose VeriSign's testbed as the winner, it seems inevitable that the Chinese and other Asian nations will start their own root. Our reviewer disagreed: "I suspect not. I suspect what will happen is that ICANN will delegate localized strings, most likely some ACE, as aliases (either NSes or DNAMEs) for the ccTLDs in the root. These aliases would point to the same nameservers as the ASCII ccTLDs. I am skeptical that ICANN would venture into the minefield of disallowing certain strings in any TLD."

VeriSign's Nominum Strategy

But VeriSign is playing a game of raw naked market power. At the URL http://VeriSign-grs.com/multilingual/strategy.html we find the following statement.

"A key element of our strategy includes adding multilingual domain and host name services to the Berkeley Internet Name Domain (BIND) software, recognized as the most widely deployed DNS server software on the Internet. VeriSign Global Registry Services has established relationships with other companies to assist with the implementation of our strategy. For example, the VeriSign Global Registry Services has provided funding for Nominum, Inc., the company that maintains the reference implementation of BIND, towards the development of a multilingual domain name version of BIND. Nominum provides BIND software to the Internet community as open source software through the Internet Software Consortium (ISC). Engaging in relationships with key DNS infrastructure providers such as Nominum, VeriSign Global Registry Services will develop a solution that provides multilingual domain name services to the greatest number of users with minimal impact on existing systems, and provide ongoing compatibility with those users who do not require multilingual domain name services."

VeriSign is saying with the previous statement that they have covered their bets in a way that the Asian nations short of starting their own root and their own DNS resolver software will find hard to match. BIND is the global de-facto standard for DNS resolver software. If VersiSign Network Solutions has paid for the development of "a multilingual domain name version of BIND," we may rest assured that much awaited BIND Version 9 will be compatible with the needs of VeriSign's multi lingual domain name implementation.

We asked David Conrad,President of the Internet Software Consortium to comment. He responded: "BIND version 9 was released on September 15. To be clear, Nominum and ISC are separate entities with different foci. ISC's charter is to provide openly available reference implementations of core Internet protocols. Nominum is a Silicon Valley startup.

ISC has explicitly stated on numerous occasions that multi-lingual support will NOT be included in the ISC release of BIND until the IETF reaches some level of consensus on how multi-lingual domain names should be implemented. ISC does, however, include contributed code in the BIND releases, some of which implement multi-lingual enhancements. To date, Verisign has not contributed any such code. JPNIC has (their mdnkit). The work Nominum does for Verisgn may or may not be contributed to ISC, depending on numerous factors."

The question it seems becomes one of whether Nominum will be able to make BIND compatible with Asian approaches designed to compete with VeriSign's. We doubt it. David Conrad commented: "It depends on which Asian approach you're talking about."

Klensin Drafts Revamp DNS?

In the face of the emergence of these apparently intractable problems, John Klensin, the new chair of the IAB has authored 3 extremely important internet drafts.

On December 13 we have http URL:draft-klensin-i18n-newclass-00.txt
Title: Internationalizing the DNS — A New Class < http://search.ietf.org/internet-drafts/draft-klensin-i18n-newclass-00.txt>

Klensin writes: "This document explores the question of what the DNS, and "DNS names" would have looked like had we been designing it today, with multilingual usage as a priority but with current technology and existing standards available, and proposes a way to get there."

"The proposal is radical in the sense that it implies a major restructuring of DNS usage and, indeed, of the Internet, to make the DNS seamlessly capable of working with multinational character sets. Such a restructuring is, and should be, quite frightening. It is worth considering only if the long-term risks and problems of other proposals are severe enough to justify a radical approach. It is the working hypothesis of this document that they are. At a relatively technical level, this would require changing every DNS resolver and server, and application that accesses either, on the Internet that wished to use non-ASCII names." Conrad: "Note that you'll most likely have to do this anyway — too many programs make bad assumptions about what data is returned from a DNS query. "

"Legacy (unconverted) systems would be at a significant disadvantage in referencing new names (some of which might use only a subset of ASCII characters but might still not

be registered in the older Class), just as legacy systems were during the transition between Hostnames and the DNS. There are also a number of problems, such as the weaknesses of the DNS as a directory system, which it does not solve (see section 3.6, below)." [SNIP]

"A "new class" proposal would obviously not be easy to deploy, but, realistically, neither are any of the other ideas if the definition of deployment involves users having access to Internet names drawn from a broad range of languages. It would cleanly separate "international character set" name spaces from the "ASCII" one — i.e., "old" clients and systems would never see the non-ASCII types. In the international name space, English, and the character code points used to represent it, becomes just one of many such languages and their corresponding character code points. It might even let us fix a few other things along the line, as long as they were sufficiently straightforward to not create significant delays. E.g., there are several RR types in the current Class that are either obsolete or have never been widely used, and we might be able to eliminate them by not carrying them forward." Conrad: "The only advantage of a new class is that, for the most part, it is relatively straight forward to implement. If you start mucking about with the underlying protocol, you might as well go and implement DNSng."

While other transition models are possible, the cleanest one would be to conclude that the new Class was intended, over time, to simply obsolete and replace Class=IN. If registrations in Class=IN were transferred into (or explicitly referenced from) the new class (or a "search rule" system was employed), then the transition model would be very similar to that of the hosttable-> DNS transition. In particular, "old" clients and systems would see a smaller and smaller fraction of the Internet until they converted and we would expect some user-level tools to arise to work around slow conversions. [SNIP]

If we are really serious about converting the Internet to support applications in all languages (which is ultimately the assumption underlying this document), then the answer may be more clear: the overhead of dealing with the UCS to ASCII interface ought to fall on those applications, as an intermediate step until the protocols themselves can be upgraded. In other words, we would anticipate a four-stage conversion process for those applications:

(i) Completely legacy (non-updated) code would continue to reference Class=IN (no other option is possible).

(ii) Applications code would be upgraded to make QClass=UC inquiries and to represent the UCS codes for their databases and presentations in some ASCII-compatible form compatible with their protocol definitions.

(iii) The protocols would be upgraded to international norms and usage.

(iv) The applications code would be changed to conform to the new protocols, eliminating the workaround of stage (ii).

[SNIP] 5.1 The root server question The design of the DNS is such that there is no inherent reason why root servers for a new class would need to be the same as those for IN. However, the same considerations for root server selection that apply to the Class=IN root [rfc2870] would presumably apply to the Class=UC root as well. There would be several other administrative and operational advantages for keeping the root servers the same —or at least co-locating them— as long as loads and similar factors permitted."

By calling for a universal class Klensin is in effect calling a new root into which the old ASCII based root would be folded as a subset of the new 'order'. Over time every DNS resolver would be obsoleted and replaced. It would be a bit like rebooting the Internet. But so convoluted have the DNS wars gown that at the very apex of power the thought has suddenly become one of sweeping everything aside and starting afresh.

On November 10th and December 13 Klensin published two further drafts: http URL:draft-klensin-1591-reflections-02.txt Title: Reflections on the DNS, RFC 1591, and Categories of Domains and http URL:draft-klensin-dns-role-00.txt Title: Role of the Domain Name System Klensin has thrown down the gauntlet to solve an apparently intractible problem. Observers believe that this fresh set of ideas may unleash waves of forward motion.

Klensin noted that: "A mailing list has been initiated for discussion of this draft, its successors, and closely-related issues at ietf-i18n-dns-newclass@imc.org. To subscribe to the mailing list, send a message to ietf-i18n-dns-newclass-request@imc.org with the single word "subscribe" (without the quotes) in the body of the message. To unsubscribe from the list, use that same address with the single word "unsubscribe" in the body of the message. Issues related to the relationship of the model proposed here to general issues of multilingual access to the DNS should be raised in the IETF IDN WG working group, see http://www.ietf.org/html.charters/idn-charter.html." We predict that this mail list could be one of the most important lists in a long long time. Those to whom these issues are important are well advised to join.

An Uncertain Future

The problems of internationalized domain names look to be intractable. If so the new class in Klensin becomes UTF8 and creates and new root into which the old one is folded. Creating a new class would offer the opportunity for old names to migrate new internationalized root. For if it is to work there would have to be a new resolver that would know that if no information was found in the old root then it should look at the new. Since Bind 9 is to be found in almost every Unix program known, a critical question is whether Paul Vixie will build in to BIND 9 the ability to handle both roots? Conrad: "Paul was not involved in the implementation of BIND 9. BIND 9 (like all other versions of BIND) already supports multiple classes (Hesiod and Chaosnet being the two other than IN supported by BIND)."

Another issue is whether ICANN would try to exert control over the new root. Thinking is that the IETF-ITU designers may have had enough to ICANN and may want to keep it ICANN free. Don Heath has resigned the ISOC Presidency effective March 2001. He would seem to be a "shoo-in" for the new ICANN president. Such events would pose an interesting dilemma for the new Bush administration. It is quite possible that Bush'es men may be willing to let the root go. If they were not, such action would likely drive the Asians to establish their own break away root. It seems to us that there is way too much emphasis on the necessity of a single root and the implied magic of having everyone on the face of the earth universally reachable by a single computer network. The internet we still believe will be controlled from the edges and not from the center.

What we are witnessing here is a very fundamental shifting of the tectonic plates currently supporting the Internet.

NSI is announcing to the Chinese that they will sell use to them the use of words from their own language. The Chinese have said this is completely unacceptable. In such a context, a depoliticized open standard such as Klensin is proposing will generate much interest by the Chinese. As Ted Byfield puts it: Klensin is saying with his RFCs that the problem is purely technical and totally open. Whereas the minute you say that the DNS solution is not open but is bounded by parameters that in effect make it both a closed system and a proprietary one in a commercial context, you are in effect closing it and selling it and making it both a political and perhaps an intractable issue.

Instant Messaging Coordination of People and Devices Becomes Standards Track High Priority

Serves as an Enabler for Many New Applications, New Uses of Bandwidth and Intelligent User Agents

Editor's Note: Since 1996 Henning **Schulzrinne** has been a faculty member at Columbia University in the Department of Computer Sciences and Electrical Engineering. From 1994 to mid 1996 he was previously at the German research institution called GMD Fokers in Berlin. In 1992 and 93 he worked at AT&T Bell Labs which from a legal perspective is now part of Lucent. He still maintains an office there and gets some of his students from there. His current interest is Internet protocols, multimedia and performance evaluation. He devotes most of this time at Columbia to running the Columbia University Internet Real Time Laboratory where he directs about a dozen a Master's and Ph.D. students. About 10 of his students are doctoral students. Some work there full-time and some part-time while employed by other research laboratories such as IBM. IBM, for example, pays their salary and their tuition while academically they're being supervised by Professor **Schulzrinne**. He is also a member of the Internet Architecture Board. The URL for the Internet Real Time Laboratory web page is http://www.cs.columbia.edu/~hgs/research/IRT/

COOK Report: Due in part to the instant messaging working group reorganization at the Pittsburgh IETF in August of 2000, instant messaging has acquired a bit of a buzz lately. I'm beginning to understand that it is a bit more than a significant service being fought over by Microsoft and AOL. Would you begin by explaining what instant messaging is all about from your point of view?

Schulzrinne: When one talks about instant messaging, one is really conflating two separate functions. The first is sometimes referred to as "buddy list" or presence, which is the ability to tell if a person happens to be currently on line. This is primarily relevant for people who use dialog services. It may be used to determine if someone is available to enter into real-time chat. The second and separate application is the message capability. This enables you to send mostly text messages to one or more people whom you have enabled to receive the message while they are online.

COOK Report: As you type it appears in real time on the recipient's screen?

Schulzrinne: Almost. Under the old UNIX talk program, as I typed in each character, that character would appear on my recipient's screen in real time. Instant messaging is indeed a message service in that I type my entire message and only when I hit return does the message appear almost instantaneously on my recipient's screen.

COOK Report: Therefore I don't suffer the distraction of seeing your message gradually appear on my screen — especially if you're a slow typist?

Schulzrinne: Exactly. There is another somewhat related service. It's a rather old service, although it's still used in quite a bit. It is Internet Relay Chat (IRC). For many years IRC has been an open Internet protocol. It is used primarily in so-called Internet "chat rooms" where groups of people converse with each other. The primary application which makes instant messaging rather different from IRC is that instant messaging tends to be only from one person to another person where the people involved already know each other. So it's a bit different than email and it is not generally used to send a random note to someone from whom you may have gotten a business card.

Why the Interest in Instant Messaging

Instant Messaging seems to be very popular within corporations as a quickie coordination mechanism. For example when I was at a conference in Israel and my student and collaborator Jonathan Rosenberg was in Australia, I had a paper to present that afternoon. We had some last-minute things to do. The paper was in New York. He was in Austria and I was in Israel. By logging into the New York machine, we were able to use an instant messaging application to figure out what to change and how to fix things at the very last minute.

Instant messaging use in corporations now tends to replace a lot of one line e-mail messages. One disadvantage, or, depending on how you look at it, one advantage of the instant messaging exchanges is that no permanent record or archive is left. In this sense instant messaging is also more informal. It certainly is not suitable for message attachments. Also, it only works as long as the person to whom you are sending remains online. You could spend five minutes composing an answer and should, in the meantime, that person log off, your carefully composed answer vanishes.

COOK Report: Other than being very popular on America Online and being of very intense interest to Microsoft, how would you characterize the reasons for strong interest in the development of protocols and standards elsewhere in the industry? What is driving all this?

Schulzrinne: Instant Messaging is apparently seen as a means of communication which is very important to the people who use it. From a consumer services perspective it likely ranks behind the Web and email as the third major reason by which people justify their Internet usage. Standardization matters because you now have a problem that there are isolated communities where people who use the Yahoo instant messaging client cannot send messages to people who use either the Microsoft or the AOL client. AOL is appears to have 90 percent of instant messaging users while the remaining 10% are split between Microsoft, Yahoo, ICQ and Tribal Voice.

While this lack of standardization is not now especially daunting in the U.S., it will also begin to get much more important once more advanced wireless services come into play. You need to be aware of how popular the Short Messaging Service (SMS) is in Europe on the GSM mobile phone. Moreover this popularity exists despite the fact that you have to type in your message with the number keys. On completion these messages are sent instantaneously to another mobile phone with bridges to email also available.

I would not be surprised if these devices become more popular in the U.S. as well and I expect that there will be bridges — even on the just chatting side between mobile devices where composing a quickie one line instant message might be O.K. but where composing a formal email would be overkill. In Europe the Short Messaging Service sends billions and billions of messages — in part because doing so is less disruptive than making a mobile phone call especially since you don't have to worry about reaching the person by voice in the midst of

inappropriate circumstances.

it is currently important as a consumer-to-consumer application and to some secondary extent it is starting to be used as a consumer-to-business application. I have heard of a number of technical-support places and a number of consumer mail-order operations which allow users to chat with their sales agents. For example on a Web site you might click on a link and have an instant messaging client pop-up. From the point of view of the provider of the consumer goods this it is more efficient than voice communications because they expect their customer service agents to handle about eight simultaneous chat sessions. To sum things up Instant Messaging is one of those strange sort of applications which is actually more popular with the nontechnical crowd than with your "techies." This has had some undesirable consequences because it meant that this issue did not appear on the radar screens of the IETF folks until very recently.

Possible Uses

COOK Report: Supposing I am the edge chief information officer of a Fortune 500 company and I want to offer Instant Messaging to my company's employees, how do I do this? Do I have to license it from AOL?

Schulzrinne: There seem to be different models out there. One model is that some corporations might just use the ICQ or Yahoo server. This has the implication that you better trust the server to be reasonably secure so that whatever private internal corporate information you exchange does not end up in the wrong place. This is also actually how many of the smaller companies tend to make money. From what I can tell their public chat servers are really an advertising shingle that they hang out there on the public Web and say: " oh by the way if you like this and want to run one internally on your Corporate Intranet, here is where you can send check for the server and we will set this up internally for you." Also it looks to me that Instant Messaging in corporations is not a service that is decided by the CIO. Instead it is something that employees tend to pick up on their own. Someone with in an engineering group says: "this is neat." I have used it at home, and now I can go back and tell all my colleagues about it and my group will begin to use it at work . This seems to be something that is more grassroots, bottom-up driven than something that arrives as the result of a corporate decree.

COOK Report: I've heard the belief expressed that geographically distributed corporate users are likely to be interested in Instant Messaging as a means of reaching out and pinging their colleagues to coordinate the start up of other bandwidth intensive activities such as video and or audio conferencing.

Schulzrinne: For sure this sort of thing will be an activity. This is why I was careful to separate my description into the part that signals presence or availability and the message part. The presence part simply indicates whether or not you're available without really doing anything. It uses heuristics to determine whether you're logged in and whether you have typed anything or not for a specified number of minutes.

Instant Messaging is primarily a first order mechanism. In other words you use it to set off other events. This explains where the interest of those of us on the multimedia side came in. What you can do is set up a number of simultaneous AOL, or whatever type, of presence sessions or messaging sessions. When your group is present, you can start up a completely unrelated application such as a voice-over IP conference call.

COOK Report: Suppose you were working in your Columbia University lab and you had three collaborators in three different parts United States. Could use the system so that the moment all four of you were on line simultaneously it would initiate a voice-over IP conference call between you? It would function like a trip wire to establish your meeting with out any of you having to be consciously aware of the need to do so.

Schulzrinne: I am not aware of software that currently has this capability. But this is exactly one of the services that we are putting into the tool which we are building here at the university as a prototype. Namely as soon as some specified condition is met, the system will call up these specified people. Currently these presence indicators tend to be approximations of reality. For example they generally don't know that someone else might be physically in my office and therefore I may not want to session started automatically. So you need a bit more mediation. The general idea, however, is exactly that. Once you have this kind of integration, it becomes much easier to set up multi-party telephone calls based on the simultaneous presence of a defined set of people.

COOK Report: Are you working on additional capabilities?

Schulzrinne: What we have currently is instant messaging as a text based application. Even if you don't think of it as a call model and view it more as a social kind of thing, it is replaceable by more multimedia oriented means of communication. For example instead of the text based messaging service, you could set up something that functioned in the background as a kind of walky talky service. This wouldn't be anything like a formal telephone call. Rather it would be a means of you saying something every now and then like your transmission on a CB radio to give a listener some idea of what you're doing, and some sense of where the listener could respond to you.

COOK Report: How about having an instant system turn on Web cams in its users offices?

Schulzrinne: It's possible to do it. But it does raise privacy concerns. For example if you don't want to be able to turn on someone else's camera-at least not without the other person's knowledge of what you are doing.

COOK Report: Could you have a menu that in effect of functions as a protocol where users could specify under what circumstances or at what times they are willing to have their various devices activated?

Schulzrinne: The instant messaging protocol is certainly an enabler for most of these kinds of functionalities that are actually being differentiated on the end user's system. The protocols that are being discussed probably all facilitate this to some extent or other. We believe that when people discover the kind of integration available, they will do more than just send text messages around. For example in Sweden there is a service available for playing joint games like Quake and things of that nature when two would-be players find each other on-line.

COOK Report: In a corporate situation could you have it start a white board? Such that a remote person would also have a white board on his or her screen and could draw something and have it appear on the screen of the originator as well as on the screens of everyone else who had joined the session?

Schulzrinne: Absolutely you could start a white board or some other corporate specific applications such as a modeling program or a project planning application. Whatever application you came up with would be able to use an instant messaging system as a set approach for initiating and managing the interaction of people who need to participate at one time in a given project. It is true that the current expectation of instant messaging is that it tends to be text only and is of somewhat limited applicability. In reality even in the area of personal communication a fair amount of extension beyond just text is possible.

Messaging for Notification of Events

There is also another dimension which I think is interesting to look at. The presence of people is just a special case of a more generic problem — namely that of notifying someone else of the occurrence of an

event. Presence is really a proper subset of what some people call event notification.

COOK Report: And some examples of this would be?

Schulzrinne: There are a number of them. For example I'm working with some people at Telcordia who are interested in using SIP to control things like home appliances: toasters, or light bulbs, or heaters for example. They think of the events in this environment being something like the temperature dropping below 50 degrees Fahrenheit.

COOK Report: In other words I could use this event based spinoff of instant messaging to remotely control devices in my home and be notified when the furnace went on or the front door opened — having that notification appear on the screen of the CRT in my office?

Schulzrinne: Exactly. This is one example and a home-based environment where people do not generate events as to whether they're offline or online, but sensors notify you of events. What you have operating here is a very powerful mechanism for wide area access to your home appliances. These events could also be thought of in the sense of alarms. You might have a sensor that would turn on a Web cam in the event of movement and send a specified alarm to a remote computer screen. In an industrial setting an event that would trigger an instant message might be an increase in the pressure in a boiler beyond a specified level.

If you take the earlier instances of someone's presence and consider it as but another category in an overall catalog of events that can be signaled to a computer on a network as they occur, you end up being able to reuse a very significant amount of software infrastructure. You avoid having to invent say a new system for industrial alarms, a new system for a home events, a new system for finding game buddy's and the new system for finalizing a text drafts.

What I see here is a powerful new enablers. We have something here that fills a gap in the kind of communication capabilities that the Internet has heretofore been unable to offer. We have had instantaneous retrieval in web pages. We have had asynchronous messaging with no immediate feedback in electronic mail. But we have not had this type of more rapid asynchronous communication with time scales measured in seconds rather than in minutes. It is much more of a push mechanism than web page retrieval. And therefore it fills an important ecological niche in our overall group of protocols.

AOL, Microsoft and Instant Messaging

COOK Report: How does the dispute between America Online and Microsoft effect what we've been talking about a so far with instant messaging? To what extent, if any, has it been responsible for the development of an IETF interest in the standardization of instant messaging?

Schulzrinne: The emergence of instant messaging as one of the primary Internet applications has certainly gotten a lot of people interested. This is very likely the largest Internet application which is strictly proprietary. The web and email are certainly reasonably standardized. Instant messaging suddenly came out of nowhere because its community of users was somewhat different than the technical community out of which the Internet sprung.

COOK Report: Has AOL not licensed the use of their messaging software to anyone?

Schulzrinne: I am not sure what their licensing policy is. They control it rather strictly because I believe it is an important ingredient in maintaining their customer base. You can actually get the software for free and I believe that under nondisclosure you can license the details of it. However AOL places strict restrictions on who can offer their instant messaging software. The result is that in not everyone can just walk up and say "I want to build an AOL compatible system. "

The protocol apparently is not that difficult because Microsoft, for example, did not license the technology from AOL but simply reverse engineered it. In doing so Microsoft made their product look like the AOL software. They observed how it functioned on AOL. (There didn't seem to be a lot of encryption going on). In one sense it could be said they didn't really have to reverse engineer the code. Instead they observed how it functioned and built software that emulated its functions.

AOL, for its part, then looked at the Microsoft emulation and tried to figure out where it was less than perfect. These minor differences didn't really have functional implications, for example, I'm guessing now that there might have been an extra space somewhere or a small difference in timing. Nevretheless, they used the differences to examine a messaging package and determine that it was a Microsoft emulation and not a genuine AOL original. When it made such a determination, it then acted to turn off access to from the Microsoft product to it's own software.

However, from the AOL perspective this is not really workable in the long term because eventually they will find something that on the wire looks exactly like their product. As long as they are not using encryption they will be unable to prevent such a product from interconnecting with their own. Still as long as AOL does do explicit licensing, it gives them some control as to who uses it. AOL is able to say their software is it the " real thing" and contrast it with what they present as an inferior imitation.

I don't know what AOL's revenue model for it is since it is available to people not on AOL and you don't have to be an AOL customer to use it and, as far as I know, it is not advertising driven. Moreover, it is easily available since Netscape, which is now part of AOL, comes with it.

COOK Report: So was it AOL that started it up in the form that it has taken now and got the critical mass going as everyone jumped on the bandwagon?

Schulzrinne: Yes. Except that I don't believe they invented the service. It was always a service that existed on the traditional BBS systems. Now AOL comes out of this heritage. When they went more Internet based, it could be argued that they took instant messaging along with them. It seems to have come very much out of the nonacademic and non technical community.

The IETF Standards Process

COOK Report: What then really happened within about the past year the got it moved into the IETF?

Schulzrinne: A working group was formed within the IMPP group to look at a standards-based protocol for this. The group was so diverse, because in this field people come from very different backgrounds, that effectively no technical progress was made. Realizing this, at the Adelaide IETF we essentially set up a design competition where the first working group and others were or challenged essentially to put up or shut up.

COOK Report: But what else happened? I understand there were two groups set up one led by Marshall Rose and Dave Crocker using Rose's "blocks" proposal and the other led by yourself and Christian Huitema from Microsoft and involving Sip.

Schulzrinne: That is correct. But unfortunately it's even more messy than that. The working group did produce a requirements document that was kind of a motherhood and apple pie compilation of what instant messaging systems should do. That was RFC 2779.

COOK Report: How exactly did you wind up with the design competition? Was this Scott Bradner's idea?

Schulzrinne: Basically the outcome was that the chairs of the working group were informed that their effort wasn't going anywhere. They were told: "Let's see if we can make progress by having a design competition." With the deadline set this channeled new blood into the effort. There was a flurry of activity. By the date of the deadline we had a least seven or eight proposals submitted with it some form of documentation. AOL was among them and submitted an over all architectural wish list. However, they never submitted an actual set of protocol specifications. The other groups submitted something close to complete specifications.

COOK Report: Was there a late June deadline for submitting the proposals with enough time to sort out the ramifications by the Pittsburgh IETF?

Schulzrinne: Yes. A small set of people was chosen to evaluate the proposals. What they came up with was effectively we really don't like any of them but we can't make a choice either. Really a non verdict. After some discussion in order to try to make things more manageable the group found that the proposals fell into three categories: the XPP or blocks-based proposal and then the SIP based proposal which a number of us including Christian Huitema worked on. Within this one, the primary work had started much earlier with a proposal within the SIP working group which Jonathan Rosenberg and I had been working on for some time. Finally there was a whole set of proposals which was called "group two", and which was characterized by a more limited single set of functions for text based presence indication and instant messaging. Since there was nothing basically incompatible in this group of proposals, they were combined into one. As a result, after this winnowing exercise there were three proposals left available for consideration.

I believe that all three proposals would be able to deliver a service that is roughly equivalent to the existing chat and other instant message services that we have today. The three groups have somewhat different design assumptions or starting point. The first one, the blocks proposal, starts out with a lower level infrastructure which they see as a generic mechanism for a number of as-yet-unspecified messaging type applications. They saw instant messaging and presence as a group that they could use for the evaluation of their model. My perception is that this is an attempt to provide a generalized messaging framework which basically allows you to send messages in a way which is roughly (and at a very high level) similar to how email works — just with more expeditious delivery and with no notion that a human is necessarily at the other end.

The proposal is an attempt to take the same model of service, namely a message that is delivered to one or more recipants asynchronously, and apply it to a set of applications rather than just layering it on top of email. The approach is interested in taking the messaging services and making them available as an integral part of our infrastructure to other applications. The framework itself is very much still being defined and far from being finished.

From what I can tell there are a number of details as well as a number of more fundamental issues being worked upon simultaneously. The working group for XPP meant only for the second time in Pittsburgh. The underlying infrastructure is still very much being defined as to whether it will run on TCP, which is the initial target, or to see whether it should also run on UDP. In that sense much of the underlying infrastructure still has scaffolding on it and I don't think anyone would claim that it is either finished, ready, or tested. They still view it largely as a standing application and are not too concerned about integrating it with anything else. They say in effect we have this emerging blocks tool set which may be useful and let's see what things look like when we apply it to instant messaging.

The SIP Instant Messaging Approach

COOK Report: What then is the SIP approach being taken by yourself and Christian Huitema?

Schulzrinne: Our approach came from a notion that we see messaging in the more generic sense of event notification and that we see a signaling system consisting of a push part and a pull part if you like. The push part comes in where you call me on the phone to find out if I'm available. The phone rings and either I pick up or I don't. The pull apart would be when I tell you, just in case you wanted to know, whether I'm available to talk or not. These two are mirror images of each other and thus it makes sense for them to be provided within a similar overall signaling framework that includes the ability to reach end systems.

COOK Report: Presumably this is where SIP comes in?

Schulzrinne: Exactly. This is where SIP comes in because it already provides functionality which we consider useful in this context.

COOK Report: So the intent would be to bring new functionality in where the old way of doing it would be merely to find out whether someone was on line or not? And in contrast, what you are developing goes well beyond the issue of whether your are on line or not. If it finds it that you are not, it will reach out and talk to your machine in order to try to get more information as to where you are?

Schulzrinne: Yes. One of the abilities that SIP has is this finding ability. And unlike traditional messaging protocol's it is designed to have entities in the network which are needed to locate you in the event of a telephone call. In locating you it will ask that you send me a message as to when you will be available. The person asking to be notified (the subscriber) and the person or object doing the notifying (the notifier) can effectively move around. So you would not subscribe effectively to a piece of hardware or an IP address. Rather you subscribe to a person. As long as the network knows through some registration system where the person might be, it will find you, that is as long as you are willing to allow yourself to be found.

COOK Report: So this would allow potential users to do what? Register the machines that they use? Put in a schedule? When you say it can reach out and find someone just exactly how would it do this?

Schulzrinne: Let me give me an example. Say that you are described as cook@cookreport.com . That would be your published e-mail address for example. In SIP we use e-mail addresses so that your telephone, instant messaging presence, and e-mail identifier would, in almost all cases, be the same. The idea here would be that you would register to receive calls or notifications or to be subscribed to mail, messaging and telephone. All of this works pretty much through the same infrastructure. That is regardless of which ISP you happen to be using , I would register with my designated externally visible address — in your case cook @ cookreport.com. Your registration would be held on a server that would be the definitive repository of your whereabouts. Doing this would allow you to be in several places logically speaking at once. For example you could both have your home phone and your office phone registered with your identifier. When you had an incoming call the server would see that both office phone and home phone rang simultaneously and whichever phone was picked up first would receive the actual call.

COOK Report: So might this be something like a GUI interface on a Web browser where a user might fill in, a month in advance, what he thinks is scheduled to be?

Call Processing Language

Schulzrinne: Almost. This would be handled not by our SIP instant messaging client but rather by our Call Processing Language (CPL) which would enable you to do precisely that sort of thing. This works as a language that you up load to the server. Although as a user you would not know you were doing this. It would be exported from your calendar and would pick up from the calendar the necessary e-mail addresses and telephone numbers. It would work so that if you are staying with a friend, you could put in your calendar an entry that: on Saturday I am going to be at such and such a place. You would give your friend's phone number in the IP address of the phone. Using the call processing language you would apply the results to the server. Then when a request comes into server to talk to cook@cookreport it would it check its listing for you and say: oh today he can be found at such and such IP address. It would forward the incoming call to your friend's house and, just in case you weren't there, it would ring your home phone as well. Thus either manually or by some logic that is generated by some explicit specification or derived from things like calendars, you are routing incoming calls, or the incoming messages in general to a variety of places. In the case of voice communication, if there are gateways available, this can be regular telephone numbers or advanced telephone devices.

So the power of using SIP not for a just an "I call you up mode" but also for an "I notify you when I'm available mode," and an "I want to find out when you are available in order to subscribe to" mode means that all of these messages are handled by the same infrastructure.

COOK Report: What do you mean by the same infrastructure?

Schulzrinne: We mean the same set of software or so-called proxy servers. You can think of those as being cousins of your SMTP server. They are similar in where they are likely to be installed. For example at ISPs and at large institutions and in some departments of large corporations and so on. They have a similar functionality in that they do a form of generic message routing. In a sense they're more dynamic than a mail server where your mail ends up and forces you to go to the server to pick it up from wherever you might be. The idea in our work on instant messaging using SIP is that the message routing can be much more dynamic and follow you around in the sense that you can tell messages to go dynamically to your and system rather than waiting to be picked up. With e-mail you pick up whenever you want to.

The work that we are doing is designed for messages that are meant to be picked up in close to real time in the sense that you're actually there when it happens. You want to get the phone call now and not be told about an hour after it comes in. The model is related to email but functionality and requirements are somewhat different. The speed of delivery is measured in hundreds of milliseconds instead of minutes. It is instant with the idea that the other side is already there rather than asynchronous and that I send a message and you may pick it up two minutes or two days later.

COOK Report: So you have some functionality in here that would be very attractive to people in commercial environments and, because not everyone uses the same hardware and software, standardization becomes important if this technology is to be universally usable.

Infrastructure with a Billion Users

Schulzrinne: That's true and we're also interested in the integration possibilities since SIP is likely to be used as the system for call setup (finding a person) in the third generation wireless systems. We see this as a central functionality which we want to leverage across related activities (call setup, subscription, and modifications). We see having a central infrastructure as being much easier to manage and provision than having separate infrastructures for different tasks. This is where we're coming from. We said we will have a very large set of infrastructures, especially, if as we expect, the next generation wireless systems will all use SIP internally, we will have an infrastructure with a billion users on it. We might as well leverage this infrastructure since the proxies do not have to be upgraded in all in order to have this functionality.

Consequently a proxy does not even have to be aware that instant messaging or presence is going on. Therefore a proxy built today will be a perfectly capable router of subscriptions and notifications regardless of what happens to its details in the future. To us, this capability opens up an avenue for integration which hopefully will lead to lots of interesting new services.

COOK Report: A question about the Microsoft environment. I understand it has instant messaging and something called Net Meeting as well that is used for setting up a voice-over IP sessions. I understand that Christian Huitema moved from Telcordia to Microsoft about six or nine months with the objective of taking control of Microsoft's Net Meeting and Instant Messaging and getting them ready to be replaced by what ever comes out of standardization effort. Are you willing to comment on this?

Schulzrinne: While I am obviously not a party to internal Microsoft product discussions, I can say that indeed Christian moved to Microsoft with that topic area as part of his portfolio. And Microsoft has publicly stated at a Voice on the Net conference that took place earlier this summer that it will be supporting SIP in various different ways without however making any specific product announcements.

COOK Report: If these capabilities are being developed, is it reasonable conclusion that Microsoft would benefit by incorporating them?

Schulzrinne: Yes. There are a number of companies going in that direction. There is a company called "Hear me" which is, at this point, engaging more in the traditional voice-chat environment. It is implementing a SIP-based non proprietary mechanism. So SIP is clearly happening not just in it the voice-over IP area or in the traditional conferencing area that is encompassed by Net Meeting, but it is also happening in the voice-chat area which is a somewhat separate small sub area that is a more informal, more walkie talky style if you like, kind of application. This is definitely occurring and one can only imagine that major software vendors would have an interest in participating in that. Everyone benefits in the end by not having isolated communities and having users have to download half a dozen different applications to participate, even if the applications are free. Benefits will come not by trying to keep the pipe sizes different but by adding services to a uniform " pipe" or standard.

Progress of the Standards Process

COOK Report: If you're looking at the future, can you give me an idea of what your schedule looks like to conclude your activities? Do you think things will go smoothly, or are there any significant areas where you have it important problems to surmount?

Schulzrinne: I believe that we have now a basic infrastructure in place to provide capable instant messaging services. Areas of future work include extending to non personal presence for home appliances or industrial applications the type of the events which we have just discussed. But this is really not so much of protocol issue. Rather it is more a case of deciding that these are going to be expressed as XML GTLDs which are going to be specific to particular domain of application. For example chemi-

cal factories will have different needs than most households. They will probably have different descriptions of their states and events, but they can be carried in the same protocol and therefore some amount of effort needs to be put into describing possibly nonhuman events.

There is also an effort to describe how the Call Processing Language can be extended to handle subscriptions and notifications that are more specific to the service itself. For example you might allow someone to subscribe to your presence in order to be able to ascertain it, but you may not allow that person to call you at midnight. You might therefore want to be able to distinguish between two such actions. This is the type of thing that the Call Processing Language is supposed be doing and that still needs to be worked out.

The basic instant messaging mechanism is complete and are a number of implementation methods on going. At the last bake off there were two independent implementations. We are working in the third one. I am sure there are others who are working on it without necessarily telling everyone about it. We found we can build workable systems and, as always, there will be extensions to these systems as people find out what new services are wanted.

COOK Report: So is it accurate to say now that you have a completed specification and what you are doing is really working on implementations?

Schulzrinne: Yes. I would say there's not really much of an issue about a maturity of a specification. We really don't have a whole lot of open issues. It is really more a question of deciding where this whole effort is on going inside the IETF. SIP based instant messaging depends, as to whether it can go to proposed standard and first implementation step, on what actions the IETF area directors and the IESG decides internally to take. One suggestion has been to decide whether it makes sense to form separate working groups to basically finalize two or three of the proposals. And then basically let the market decide which one to use.

Now there has been a joint draft put out by the representatives of each one of these proposals explaining how interworking between these proposals could work in the marketplace. Even under the worst case assumption that all three proposals progress, we will be a much better situation than we are today. This will be true because, even though people use different protocols which is somewhat inconvenient, they still can use, across the protocol communities, at least the common intersection of the services between the three which will be basic instant messaging and presence. So this is work which still needs to be fleshed out in some of its nitty gritty details. This is one possible outcome.

Therefore it is a little bit early to speak about the standardization schedule for any one proposal because this schedule depends a bit more on the macro level discussions as to where things are supposed to be going — all of which seems to be less than clear at this point.

What is happening is that the leadership of the Instant Messaging and Presence Protocol (IMPP) working group will likely change. This was the original chartered working group. Different leadership for this group will hopefully lead to a resolution on which way to progress. There are a number of conceivable avenues, one of which is to have the Area Director deputize representatives and a bang together a sufficient number of heads until to the proposals give up or merge somehow. Let me say this outcome is purely hypothetical. What I see as more realistic is that since the assumptions underlying the three proposals are very different, two or three of the proposals move forward as independent working groups. What will then happen is that, after some time, it will become clear whether anyone outside the working groups actually cares about the output of the working group. In other words it is possible that one or more of the working groups will wither from actual lack of outside input. Or we might find that the working group can produce a spec in a reasonable amount of time and had people actually began to implement it. So if it turns out to be a full-fledged battle in the marketplace, it may simply be a matter of finding out who is most efficiently motivated to push the work forward.

December 2000 COOK Report

Part Five

Other Oncoming Technical Developments that Devour Bandwidth

We conclude our suvey with a group of miscellaneous interviews each of which highlights on coming changes in the intenet that will demand more and more bandwidth.

Data mining is acquiring public standards and, with the arrival of these and the XML markup language, we may expect to see the development of a Data Web where huge datasets can be remotely interogated and can even interogate each other - an activity where each application can use hundreds of megabytes per second while in progress. IPv6 will begin spotty deployment as millions of digital wireless devices start to come on line. Napster marked the debut of peer-to-peer applications where users store and share personal sound libraries on the internet. Spread spectrum wireless is beginning to power ISPs and is being adopted by scientists for envionmental data collection. Finally Akamai is tying to find the best way to manage the growth of the document web by distributing its data.

Executive Summary

Data Web pp. 208-219

A new web is emerging. The data web will likely exceed the document web in size and in its impact on Internet infrastructure. We interview Robert Grossman, CEO of Magnify, Inc. and Director of the Laboratory for Advanced Computing at the University of Illinois Chicago. Grossman has played a pioneering role in the use of high performance computer networks to assist scientists in their analysis of extremely large data sets. He has built a layered view of how data mining - a process of data analysis and real time decision making - could be carried out over the Internet.

Many businesses have extensive data sets showing information about their customers including their customer purchasing history. Grossman explains his role in catalyzing the Data Mining Group which is a consortium made up of Angoss, IBM, Magnify, Microsoft, Mineit, NCR, Oracle, Salford Systems, SGI, SPSS, and Xchange. The group is made up predominantly of vendors of proprietary data mining software packages. These vendors are now joined in an effort to develop a set of open standards that should lead to much new software and to a vast increase in the amount of data mining. Furthermore, with the spread of the XML markup language that is used to display rows and columns of data on the web, it is expected that these developments will lead to the take off of a public data web. This will mean the growth of sites having publicly accessible data sets where visitors with client browsers equipped to interact with the site's data servers can retrieve data that can be manipulated as data rather than examined but not changed as is the case with an HTML page. The result will be the data web or what Grossman calls Data Space.

As Grossman explains: "From the user's perspective, Data Space works like the document web. You can use a browser to examine remote and distributed data. And you can analyze and mine it with a point and click interface. Web sites can use Data Space services such as personalization and predictive modeling to provide a site with interactions which are created on the fly for each individual visitor.

From the vendor's perspective, Data Space is also like the document web, it simply uses a richer suite of services, including services for moving data (DSTP) and real time scoring (PSUP), and specialized XML languages for working with data, including the Predictive Model Markup Language (PMML) and the Data Extraction and Transformation Language (DXML)."

Data Space uses open standards to provide the Internet infrastructure necessary to work with scientific, engineering, business, and health care data. Unlike HTTP and HTML which are designed for multi-media documents, Data Space is somewhat more complicated because you have higher expectations when you work with data than when you work with documents."

"A document you only have to read. With data you have to analyze, score and make decisions. What everyone interested in tracking and planning for the further growth and development of Internet infrastructure needs to understand is that so far the current internet barely scratches the surface of what you will be able to do with data as Data Space and similar infrastructure begins to be deployed. I'm sure that the data web will be an important driver of bandwidth over the next few years."

Grossman also explains the Terabyte Challenge <http://www.ncdm.uic.edu> which for the past four years has been used both as a test bed for the basic protocols, languages and tools for Data Space, as well as a testbed for different ways to scale data intensive applications, especially remote and distributed data intensive applications. The focus has been on developing an open infrastructure for working with large and distributed data sets. Grossman's group has developed a process of stripping that allows large data sets to interact with each other in real time at sustained bandwidth usage of more than 250 megabytes per second.

The data space transfer protocol (DSTP) is the protocol used to move data between nodes in the data web. The data extraction and transformation mark up language (DXML) describes how to clean, transform, and shape data. This is usually one of the most labor intensive tasks when working with data. Statistical models are built using statistical and data mining applications. The predictive model markup language (PMML) describes the output of such systems in an open format. Scoring is the process of using statistical models to make decisions. The Predictive Scoring and Update Protocol (PSUP) is a protocol that can be used for both on line real time scoring and updates as well as scoring in an off line batch environment.

When PMML was adopted as an open standard by the likes of IBM and other major players earlier this year, the trade press published a flurry of articles. However, our interview with Grossman represents the first article that covers the entire extent of what he is doing.

IPv6 Mobile Communications, pp. 220-226

We interview Charlie **Perkins** a Research Fellow in the Wireless Internet Mobility Group at the Nokia Communication Systems Laboratory. Perkins offers a fresh point of view on the issue of IPv6 deployment. He explains that independent nodes running IPv6 already exist and will spread. "IPv4 and IPv6 can co-exist in the same general network because they do not collide with each other. They just have to know how to address each other. For example you can have a router that routes IPv6 packets and IPv4 packets on the same network."

"The whole thing about IPv6 to begin with was to develop a protocol, deploy it, and do what IETF does well which is to get to inter-operability testing going and then to just start to build it. People want to buy solutions to the problems facing them, be they IPv4 problems or IPv6 problems. People will want to buy solutions for their IPv4 problems and for their IPv6 problems. Eventually the solutions for the IPv4 problems may become more expensive than the solutions for the IPv6 problems. This will be true in part because the IPv6 solutions that are already available will become cheaper as IPv6 grows in market share."

NAT will not suddenly disappear. Having large domains of both IPv4 and IPv6 is merely one way to partition the possibility of overall IP address space in general. In such a situation with the right kind of fire wall NAT platform, you can even mention translating IPv6 addresses into IPv4 addresses at the border of the domains so that IPv4 applications can in effect be tricked into believing that what is going on is only an interaction between two IPv4 applications." Talking about the inordinate expense of converting and IPv4 internet into IPv6 is asking the wrong question because v6 can be meaningfully deployed in an Internet where v4 continues to function.

However the arrival of a billion cell phones over the next 18 months will force much more serious deployment of IPv6 which is the only reasonable means of doing both voice and data over a single cell phone.

According to Perkins: "We have answers for most of what we have looked at but, as we

look, we see more and more problems. For example there are a lot of problems in security and a lot of problems in quality of service. There are also a lot of problems in header compression. Also the way in which the base stations are coordinated to manage spectrum most effectively is historically not very friendly towards the IP model."

"All of these things add up to a situation where, as I mentioned before, you can employ IPv6 now. But for specific applications like voice IPv6 cannot currently match the performance of analog voice-over the air as a part of the PSTN. Now we're going to change this and believe that we will be able to equal or exceed the current capabilities of analog voice over the air as part of the PSTN." Adding mobility to the mix of necessary protocol development for IPv6 data phones complicates the technical issues involved. According to Perkins: "There are a lot of people who want to use v4. But I don't think we will ever get to global deployment of mobile IPv4 for voice-over IP. I think by the time voice-over-IP really comes into play, we will be using largely IPv6."

Of several interesting protocols being developed the most interesting is an IETF a working group called AAA (authentication, authorization and accounting). Radius only works for static objects and has some other difficulties as well. Consequently this AAA working group is building up a replacement protocol for Radius. The AAA protocol will come with features such as session measurement and Accounting. Tied in with IP SEC, AAA will do authorization and accounting for services such as mobile IP.

Comparison of v4 and v6 in Scaling, pp. 227 - 229

Yet another IETF discussion this time with interesting new information on levels of complexity of NATs and levels of address allocation.

Napster's impact, pp. 230 - 232

Napster is an application written by a 19 year old computer science student last summer. Downloadable from the web, it lets users temporarily turn their computers into servers for the purpose of swapping MP3 files. Grown hugely popular in the last several months, it accounts for a significant percentage of Internet traffic. According to university network administrators, it is clogging campus connections to the Internet. We publish edited discussion on what can be done about the problem from the CAIDA and NANOG mail lists. Port blocking has been tried without great success as student s in many cases find other ports to use. A new program called Gnutella and far more powerful than Napster is under development as well. Some folk are saying that Napster's impact on Internet traffic may approach that of the web.

Hughes on Wireless pp. 233-247

We interview Dave Hughes, principal investigator of two NSF wireless projects and owner of Old Colorado Communications on the state of the art of TCP/IP radios for wireless ISPs and for scientific environmental field research (his second major NSF project). He points out how a fall in prices and increase in capability has pushed the price of 10 Mbps radios to well under $1000 each.

He views Cisco's purchase of Aironet as a major move forward for the wireless spread spectrum industry however the impact of this will ultimately depend on how Cisco integrates Aironet products into its business line and whether it comes to see itself as a provider of connectivity solutions. Cisco has announced its own 45 megabit per second LMDS radio and is also bringing out a line of UNII band radios limited to five miles in transmission range.

These new radios can be remotely logged into and configured - something that greatly increases their utility for ISPs. ISPs meanwhile are going wireless. Breezecom claims ISP 500 customers in North America. Some 100 megabit per second radios are beginning to appear. One such is made by Proxim.

While line of sight problems are critical for these radios, Cisco is claiming to have overcome some of the drop off of communications caused by Fresnel Zone problems. Subtle physical differences found at each site can mar transmission capability. ISPs must have staff skilled in installation. The key business model is focused on connecting small business and will be increasingly focused on delivery of broadband services to residential customers who either don't want or can't get adequate DSL or cable connection.

Hughes discusses in detail the way that the E-rate increases by an order of magnitude the cost of connecting public schools to the Internet by prohibiting the schools from buying wireless equipment and requiring them instead to rent leased lines from the local exchange carrier year after year. Under the Texas subsidy, the monthly cost of the Internet connection is multiplied by 12 and to that figure is added the cost of hardware (DSU/CSU for example) needed for the phone connection. The total sum becomes the amount of subsidy for which the district is eligible. The district is then free to spend the money on the telco connection or on radios and a wireless connection. Wireless normally wins because the district after costing out the alternatives, normally finds that the cost of radios and plug in via radio to an ISP leaves them several thousand dollars left over.

At the 2.4 giga hertz range most manufacturers make radios that operate at one tenth of a watt or 100 miliwatts rather than the allowed power of one watt. They do this because [such radios can be sold in Europe, and] it saves considerable money in the cost of manufacturing. Customer don't seem to mind because if their cheaper 100 miliwatt radio won't connect they can buy an inline amplifier for $750 and increase the power to a full watt. When this has been done and under ideal line of sight conditions the radios have successfully work over distances of up to seventy miles.

Given the lack of incentive for wireline telco's to bring broadband into rural America, the FCC is has issued a notice of inquiry on the subject of a Software Defined Radio (SDR). One where smart software controls the radio - its power, its frequency spread, and other technical characteristics. Major spectrum possibilities could be achieved simply by allowing the design and use of radios that could tune themselves in accordance with the operational reality of their surroundings.

Hughes points out that one way a user who lives close to an ISP pop can help to spread the benefits of wireless technology is to ascertain whether the ISP operator with allow him to plug a radio into the ISPs pop connection. If so after a [site] sight survey to determine that radios needed are available at reasonable price, the user can install the radios and plug one into the ISP's Ethernet.

To ensure that they can inter operate with each other, radios are being built to the 802.11 standard. As shown by Apple Computer in its Macintosh Airport Base Station and Airport card architecture, it is possible for someone to spend upwards of a thousand dollars to connect to an ISP with a point to multi-point multi megabit radio and then by placing an omni directional antenna on his roof to connect and relay as many as twenty neighbors using $100 PCMCIA radios plugged into their lap tops.

Of special benefit to ISPs selling wireless connectivity is the ability in software to throttle down to agreed upon rates, the speed of the connections that radios hey supply to their customers give. Also of significance is a new Ethernet PPP protocol that means they need not assign each customer their own IP

number.

"Watch out UUNET, the frogs and the shrimp are coming - using your bandwidth" - Dave Hughes

Having laid out this general background, Hughes goes on to discuss aspects of his current NSF funded research on wireless and satellite connectivity for environmental research.

Hughes emphasizes Tachyon which has just come on line in the spring of 2000 with Concentric as its Internet providing partner as the first company to provide reasonably priced bi-directional satellite linkage to the internet. Tachyon provides a ground station that talks to the satellite for only $5,000.. It sells bi-directional, true TCP/IP, at 2 megabits down and 256k up, for $2000 a month, or 300kbps down and 64k up, for $795.

Hughes describes the National Environmental Observatory Network as part of an expanding need for environmental data collection - one that is so broad in its proposed scope that it looks as though only wireless data monitors may do an cost effective job of data gathering.

He talks about several projects in which he is working with environmental scientists whom he is surprised to find are generally unaware of the what improvement in wireless data gathering technology over the past five years make it possible to do.

The kinds of data gathering involved are quite diverse. For example the transmission from sensors the chemical composition of lake water in timed coordination with the overhead passes of a satellite. A second is the collection of light readings from a network sensors on a forest floor and the transmission of that data from each individual sensor a short distance to a data collector. The collector, in turn after perhaps encapsulating it in a tiny Linux data base, transmits it back to the research station and the Internet. A third is the capture of sound - in this case the call of the coqui frog from the rain forest of Puerto Rico. A fourth is the visual observation of shrimp transmitted in high bandwidth in real time. A possible fifth would include the use of partially buried sensors to grab, chemically analyze via tiny Linux and transmit the composition of the gasses of a prairie fire as the fire ignites the material around the sensor and passes over head. Other phenomena measured may be as diverse a earth quakes and hurricanes. From the point of view of Internet infrastructure this means yet another huge increase in bandwidth that will be generated and dumped on to backbones. Watch out UUNET, the frogs and the shrimp are coming - using your bandwidth.

Akamai, pp. 248-260

Rapid and reliable delivery of web-based content anywhere in the world has become one of the most critical issues in enabling the continued the scaling of the Internet. Web caching started out in 1996 as an attempt of many ISPs to store locally as much of the content of the web as possible. Each ISP would make its own decisions about what content to fetch and how often to do it. This system created many problems for web content providers because they had no knowledge about what was cached where, by whom and with what frequency. Furthermore, since caching distributed their content to many sites, they had no reliable way of reporting to their advertisers how many people had seen the material. It was a hit or miss system that no one was happy with and one that created a major opportunity for others to fill. A year ago Sandpiper and Akamai were the most talked about competitors. We note that since then Sandpiper has been acquired by Digital Island and has been focusing on the rapidly growing field of business-to-business e-commerce, leaving Akamai as the acknowledged leader in general content delivery.

In late 1999 Avi Freedman left his position as Vice President of Engineering for AboveNet to become Vice President of Network Architecture for Akamai. We publish a long interview where Avi explains in detail Akamai's extremely interestng business model. What Akamai does is enabled by a very significant new use for DNS that it has developed.

Akamai it has its own network of DNS servers that keep in contact with each other globally. Akamai's other servers take the web content of Akamai's customers and store it in hundreds and then thousands of copies at the edge of the network as Akamai's global network of servers continues to grow. Akamai solves the problem of the world wide wait by pushing content as close to the end user as possible.

Akamai's network of DNS servers then accomplishes a kind of global air traffic control task of communicating among themselves network traffic conditions in real time to determine which local server to send a user's request to or, in the event that regional traffic problems are interfering with local reachability, how to retrieve the data from a more distant server.

Within a site Akamai figures out what data is not constantly updated. That data is migrated to Akamai's edge servers on a regular basis. The minimum amount of data possible is pumped from the host web sites to the edges, while each edge web server is kept constantly informed of the best path to get to the fresh host data it needs. Akamai charges each web site owner for the aggregate amount of its data delivered to end users anywhere in the Internet. The table (at the end of the interview on page 12) shows how many networks receive what percent of Akamai's total aggregate of content traffic. Its intelligent overlay network of DNS servers that direct web content look up must keep very good statistics so that Akamai knows what to bill each of its customers who pay to have their web sites included in Akamai's distribution network.

Akamai has, in effect, created a virtual private overlay of the internet where, as much as possible, it keeps packets on a single network and minimizes their having to flow upstream to transit from one backbone to another (where most packet loss occurs) and them move to the downstream customers of the other backbone. This means that Akamai can go to an ISP and ask to place its servers in the its key POPs for no co-location charge and no charge for bandwidth used. Why? Because it can generally show every ISP how, with Akamai servers locally, its customers will pull far less web traffic across the ISP's backbone that they would if the ISP tried to do its own caching. Or, if the ISP just sent the packets back and forth to the content provider's central servers. In addition Akamai can demonstrate how, in return for nothing more than some co-lo space and bandwidth, the ISP will save bandwidth and give its customers better service.

Freedman also describes how Akamai must deal with the needs of its customer's central servers that are most often located at large web hosting centers at major backbone sites. In these cases he may act as an advocate for the Akamai customer in procuring if necessary some Akamai owned and operated short haul links to ensure that they can have enough burstable bandwidth to meet peak traffic periods. Given his experience at AboveNet which ran this type of operation, he is well equipped to deal with the web based, content provider, the web farm backbone operator and the large number of down stream networks where delivery oriented servers can be placed as close to customers as possible.

Akamai has taken advantage of a narrow window of opportunity to become, in contrast to the older generation vertically integrated backbones, one of a small but growing number of content distribution networks. Such a network hopes to solve problems like the peering problem for a BBN which in the summer of 1998 rebelled at granting Exodus free peering because Exodus dumped more traffic into BBN than it took out.

Three Dimensional Data Web Set To Emerge

New Protocols Enable Manipulation of Quantitative Data by Data Web Browsers -- Open Standards Likely to Give Huge Boost to Data Mining Activities -- Work Pushed By Terabyte Challenge Consortium Enables Remote Interaction of Data Sets

Editor's Note: Dr. Robert **Grossman** is chairman and CEO of Magnify Inc. Magnify provides outsourced data mining and intelligent profile services. He is also Director of the Laboratory for Advanced Computing <http://www.lac.uic.edu> at the University of Illinois Chicago. The Laboratory focuses on Internet infrastructure technologies especially involving large data sets, high-performance computing, and high-performance networking. Since 1988 he has been a Professor of Mathematics Computer Science and Statistics at the University of Illinois Chicago. For the past five years he has divided his time between Magnify and the Laboratory for Advanced Computing. From 1984 to 1988 he was a post doc at University of California Berkeley. He received his Ph.D. in computational mathematics from Princeton University in 1984. Some of his articles can be found on Two Cultures web site.

Editor's Introduction

The impact of most technology on which we report is socially and politically neutral. Fortunately! What we are caling the Data Web and what Robert **Grossman** calls Data Space is extremely powerful. The enabling of a data web with open standards will inevitably fuel data mining activities carried on now with largely proprietary software.

The data web technology itself is inherently neutral. However depending on the data to which it is applied and the use that is made of the results, the social political and personal impacts of this technology can be profound. By way of illustration we offer two examples below.

When the data mining involves human subjects, it can raise sensitive issues of privacy. For example, the News Scan Daily mail list on December 20, 2000 wrote: " In the world of struggling Internet startups, sometimes the data they've collected on their customers turns out to be their most valuable asset. "Data is worth a lot of money," says a Forrester Research analyst, and therefore "the incentive for startups to violate privacy is strong." Even the biggies -- DoubleClick and Amazon -- face alleged consumer privacy breaches, and the fear is that online companies' cavalier attitudes will create a consumer backlash. Forrester estimates that privacy fears will cost online retailers $12.4 billion in sales this year, and predicts that when Congress reconvenes next year, consumer privacy legislation will be high on the agenda, because it is a bipartisan issue that affects everyone. (InformationWeek 19 Dec 2000) <http://www.techweb.com/wire/story/TWB20001219S0010>"

These issues also can be elevated to an Orwellian point of view where they involve the privacy of an entire nation. On December 3 Dave Farber published to his Interesting pesons list the following aticle from the *Observer* in the UK. <http://www.observer.co.uk/uk_news/story/0,6903,406191,00.html> It was titled "Secret plan to spy on all British phone calls" with a by line by Kamal Ahmed, political editor

"Britain's intelligence services are seeking powers to seize all records of telephone calls, emails and internet connections made by every person living in this country. A document circulated to Home Office officials and obtained by The Observer reveals that MI5, MI6 and the police are demanding new legislation to log every phone call made in this country and store the information for seven years at a vast government-run 'data warehouse', a super computer that will hold the information."

"The secret moves, which will cost millions of pounds, were last night condemned by politicians and campaigners as a sinister expansion of 'Big Brother' state powers and a fundamental attack on the public's right to privacy."

"Last night, the Home Office admitted that it was giving the plans serious consideration. ….The document, which is classified 'restricted', says new laws are needed to allow the intelligence services, Customs and Excise and the police access to telephone and computer records of every member of the public. It suggests that the Home Office is sympathetic to the new powers, which would be used to tackle the growing problems of cybercrime, the use of computers by paedophiles to run child pornography rings, as well as terrorism and international drug trafficking."

The *COOK Report* sees no easy answer to the missuse of this technology. It exists. Complaining about it won't make it go away. One would however that in addition to promoting industry standards the Data Mining Group and others involvolved would also tackle the difficult task of educating the public and its law makers not only as to the capabilities of the technology but also as to appropriate and inappropriate uses of the results. The newly emerging data web industry needs to figure out how to apply a risks analysis to projects that see the data web tools as providing a quick fix to the dangers of anonymity of cyber criminals and terroists. An article like the just quoted one in the December 3 *Observer* should not go for more than a short time without comment from auhoratative source within the data web and data space developers.

Commoditiztion of Data Mining: Methodology and Architecture

COOK Report: What is data mining all about and how did you get involved in it?

Grossman: I first got involved when I started a project called the PASS Project. This was a data mining research project for the super conducting supercollider. It involved about 25 scientists from five universities and national laboratories. Of course when the collider was canceled, the project died. But I have always been interested in how you could do meaningful statistics on terabytes of data and make decisions based on the analysis in real time. The collider presented an interesting opportunity in how to scale up applications in three different technologies (data mining, high-performance networking, and real time detection).

We faced three challenges. The first challenge was understanding how to do statistics as data sets grew from gigabytes to terabytes and then to petabytes. The second challenge was how to provide scientists from around the world transparent access to the data using high performance computer networks. The third challenge was the real time component — namely deciding in real time what decisions to make. The collider could capture only about one out of a hundred thousand events. A decision had to be made

in real time whether to even record any particular data event.

COOK Report: Did the real time aspect of the decision making cut through everything that came after, or was this to some extent an anomaly given in the nature of the collider project?

Grossman: My interests have always combined these three areas. If you look at the PASS Project which ran from 1991 to 1994 being able to deal with large data sets in real time in a distributed manner was restricted to scientific collaborations and to defense applications and I have always been involved in both. If you move ahead to today, you will find that in the business world large data sets about customers are now common, and large corporations have distributed data about their customers. Real-time decisions about customers are now essential for the web, for call centers, and for customer relationship management (CRM) systems.

At a very basic level, my interests have remained the same for the last ten or fifteen years. But what has happened is that the computing and networking infrastructure underlying the web has grown driven by the commoditization of computers and fiber. With the commoditization of computers, there is now enough CPUs to easily do data mining even on large data sets. As far as data storage goes, while handling gigabytes used to be a challenge, terabytes are now common. Finally the technology to enable real time decisions is much better understood today. For example one of the things that we do at Magnify is to make decision in real time about customer interactions and use this information to build web pages.

What I am interested in doing is taking things that used to be quite esoteric and building them into what will become part of the common and everyday protocols of the next-generation Internet
.

COOK Report: What determines whether you study just one data set or take two data sets and compare parts of each with the other? And perhaps cause parts of each to interact with the other? Did it only recently become possible to have data sets interact with each other?

Grossman: Let me answer that by giving you some examples from some of the projects that I have been working on recently. The academic community is concerned, by and large, with a scientific and engineering data. At least at the leading research institutions, the infrastructure is two or three or four years ahead of the infrastructure at large corporations. For example at the University we've been using OC-3 networks for three or four years and have recently begun to experiment with applications using OC-12 and OC-48 networks, which is the current state of the art. With the propagation of fiber, you only now beginning to see businesses build corporate networks with OC-3 bandwidth.

One of measures that I always use is that with an OC-3 connection you can move about 10 gigabytes of data in 15 minutes. A company with a million customers may have 5 or 10 GBs of customer data. With a high performance network, instead of building a centralized data warehouse, it is possible for the first time to move data on the fly and use it for ad hoc analysis in real time situations.

Say for example that you have a site that does your Web hosting and it produces about 5 GB of data. Say you have a site somewhere else that does your call centers and that this site accumulates about 5 GB of data. Building a unified view of your customers requires analyzing both data sets, which for the first time becomes possible with an OC-3 network without first building a centralized data warehouse.

COOK Report: Under those conditions you began to think of developing a common tool set of protocols to deal with these sets of remote data?

Grossman: What we did was to build a layered view of how data analysis and real time decision making could be carried out as a vital part of the Internet infrastructure.

COOK Report: In other words issues like methodology and architecture?

Grossman: Yes. At the lowest level you move data. Then once you have moved it, you typically have to transform and clean the data.

COOK Report: What do you mean by cleaning and transforming the data?

Grossman: For example, you might have a hundred different events about a customer: recent web interactions, recent purchase, and so on. Before you decide what products to offer a customer, you generally clean and transform the data in order to compute certain derived statistical quantities. For example you count the total number of purchases, look at the most recent purchases, compute the total dollars spent, and things of that nature.

COOK Report: In other words your software must be able to formulate a set of queries to your data set that will allow you to extract the information you need?

Grossman: Yes. To put all this into a broader context, I was one of the founders of a standards group that is trying to define and standardize open source protocols for this entire set of processes. First you must move the data. This will result in a protocol known as the Data Space Transfer Protocol (DSTP). Then you need a Data Transformation Protocol. For this we are building a language right now called the Data Transformation and Extraction Markup Language (DXML). Once you have data that is clean, you can begin data mining. The output of a data mining system or data mining tool on a data set is a statistical model.

We created an XML standard for statistical models called the Predictive Model Markup Language (PMML). This standard has very wide vendor buy in. You can then take the PMML output and use it for real time decision processes. In an E-Commerce application the use might be to build a web page. In a scientific application you may use to process a data set. In the business world you may use it to build web pages or to suggest scripts in a call center. We have created another protocol called the Predictive Scoring and Update Protocol or PSUP. This will allow you to make real time decisions using PPML models.

Here is what we mean by real time. An event takes place. That event maybe someone coming to a web page. You have a PMML file in the scoring engine governing the display of the appropriate personalized content. The arrival of the person at the web page causes a profile to be pulled, the profile to be updated by the event, the new profile scored with the PMML model, and the appropriate content to be displayed.

COOK Report: By definition then this is something that can be done only for repeat visitors?

Grossman: If we are dealing with a first-time visitor you would typically start with a profile containing very little information, but the basic process is the same.

COOK Report: This must be the kind of thing that Amazon.com does when you come to browse and purchase books or other items.

Creation of an Open Infrastructure for a Data Web

Grossman: That is correct but in the case of Amazon they do it with proprietary black box systems. On the other hand we are trying to set up an open infrastructure so that anyone can do this as a routine portion of what we are calling the Data Web. Right now the web is very good for looking at and moving remote documents which can also be multimedia documents but it is not good for working with data and making real time decisions.

COOK Report: Is one of your goals to be able to compare information that you can find at one place on the Web with large amounts of similar information?

Grossman: That's right. If I want to do a point and click analysis of data available on today's Web I cannot do this in the same sense that I can retrieve a document with a point and click. We asked ourselves the question of what would be the analogies for the data web of html and http.

Here is an example to explain what we have in mind. Several months ago, the New York Times reported on an article in a scientific journal describing a relationship between El Nino and cholera. If you went to the web after you read the article, you would have found that there is no simple way to get the data and see if this is indeed such a relationship using a simple point and click browser interface.

However we already had some El Nino data from NCAR and a DSTP server in Boulder on the data web. We got some cholera data from the World Health Organization and put it on a DSTP server in Chicago. At the Supercomputing 2000 Conference in mid November we demonstrated using the data web to do a simple point and click retrieval and analysis of the El Nino and cholera data.

With our Data Space browser, <http://www.dataspaceweb.net> we can easily look for correlations between the two data sets. This does not mean that you will always immediately understand what Data Space you see. But unlike the document web, you can at least touch the data in a way that has some meaning.

COOK Report: It sounds like you are talking about giving ordinary users of the Internet tools (which indeed some users have right now) that they can use in the manner of a client and server. They will reach out to one database and say OK I wanted to talk to you in such as such a fashion. And you go to another database and set up a session and then have the two talk to each other.

Grossman: That's correct. At the Laboratory for Advanced Computing we're providing open source Data Space servers by means of which you can put your own data on the Web and we are also providing open source Data Space clients so that you can model any data from one or more Data Space servers.

COOK Report: Is it reasonable to draw any kind of analogy between the way what you just described works and Napster or Gnutella?

Grossman: Napster is peer-to-peer. What I am describing is also peer-to-peer. Napster provides infrastructure and so does this. Napster yields infrastructure for discovery and retrieval of MP3 files. This infrastructure doesn't do discovery yet which is something that will be accomplished by other tools. What we are building is a set of protocols, languages, and tools that will yield an infrastructure for (1) looking at data, for (2) doing remote and distributed data analysis, and (3) doing real time decision making. This and related protocols and languages will create what we are calling the data web. In the same way that there is a music web out there supported by Napster, we are creating a data web specifically designed for business, scientific, engineering and health care data.

The web was originally driven by people who wanted to take documents and make them available to the world. Today when scientists have data, they publish a paper. The problem is that it is extremely difficult for anyone out side the specialized community in which the data originates to make any use of it.

COOK Report: Is one reason that other people cannot "touch" it that it is not out on the web and in XML format?

Grossman: That's correct. From our experience so far, XML is a good format for meta data but not always the best format for ordinary data. The data web we are building not only supports XML, but also a variety of other data formats which may be as simple as rows and columns.

COOK Report: So what you are saying is that in the early years of the web, the placement of any data there was done without any thought of enabling people to massage or manipulate it in any way and while XML is one way to enable the manipulation of data it is by no means the only way.

The Difference Between the Data and Document Web

Grossman: That's correct. Look at it this way. It used to be that when you got a fax you were happy. Now when you receive material in fax form rather than electronically, if you actually have to work with the document, you would much prefer the electronic form given all the useful things that you can do with it that cannot be done with paper.

But think about the web. People put data there but although you can see it, you cannot sum it, average it, correlate it, nor do regression or any of a number of other analytical tasks. It is like having a fax.

COOK Report: It is like the difference between having a spread sheet that you can manipulate and printing that spreadsheet to an Adobe Acrobat PDF file that you cannot manipulate.

Grossman: Indeed. Imagine the current web data as PDF data. It is an image that you cannot do anything with without a great deal of additional work. What we have created are protocols, languages and an infrastructure so that when you see data in one of our Data Space browsers, you can easily work with it. If you want to make a decision about a customer, you can push a PPML file, you can push an event, you can pull a profile and you can make a decision in real time.

COOK Report: If you look at all the documents out on the HTML web right now, how much of the total is useable by these tools?

Grossman: Not much. Right now we are not trying to translate existing HTML data formats. We are focusing instead on the majority of data which has never been in HTML. As an example, most of the data in the deep web is of this type. It is found in flat files or in databases behind gateways. We have developed very simple ways to get the flat files into the data web. You basically just create a simple XML file describing the names and the columns and the attributes and variables. You put the resulting XML file and the ascii file containing the data itself into a directory that a Data Space server knows about and suddenly it is on the web in an entirely new way.

COOK Report: You are saying then that you can tell me how to put such data into a directory on my HTML web site that would be accessible to your tools?

Grossman: Yes, you would simply have to put a Data Space server on your website. Your data would then be part of Data Space.

COOK Report: Can your Data Space server be added to the same hard drive that my Apache web server resides on?

Grossman: Yes. tThink of it as something that can coexist with the Apache server on the same physical box. Your box probably already runs several services, http, simple mail transfer protocol (SMTP), network news transport protocol (NNTP), Napster, etc. DSTP is simply another such service. The Data Space client talks to the Data Space server through the DSTP protocol in the same way that your web browser talks to an apache server via the HTTP protocol. DSTP has more semantics than http since data has more semantics than documents, for example, data has rows and columns and ranges, etc.

How it Works

COOK Report: Is the talking done in such a way that the data would ordinarily be useable by that Apache server?

Grossman: Not quite. If you want to use the data web instead of an Apache server you would use a Data Space sever. Instead of Netscape or Internet explorer, you would use another browser. Eventually they will be merged. Right now it is rather like in the early net era when you used a different program to read news and to browse the net. Today, of course, both functions are merged into the browser. This software is now available and certainly anyone who can put up and run an Apache server can also put up a Data Space server.

COOK Report: Can you explain in more detail how this all works?

Grossman: From the user's perspective, Data Space works like the document web. You can use a browser to examine remote and distributed data. And you can analyze and mine it with a point and click interface. Web sites can use Data Space services such as personalization and predictive modeling to provide a site with interactions which are created on the fly for each individual visitor.

From the vendor's perspective, Data Space is also like the document web, it simply uses a richer suite of services, including services for moving data (DSTP) and real time scoring (PSUP), and specialized XML languages for working with data, including the Predictive Model Markup Language (PMML) and the Data Extraction and Transformation Language (DXML). Table 1 below illustrates these differences.

A user can explore Data Space directly using a Data Space browser. A Data Space browser uses the DTSP to move data and meta-data from a remote Data Space server to the local Data Space client. The data can then be cleaned, transformed and shaped either in an ad hoc fashion by the user or following the rules in a Data Extraction and Transformation Markup Language file (DXML).

Once cleaned in this fashion, the data can then be mined using one of the several commercial data mining or statistical applications available. This in turn produces a statistical model expressed in the Predictive Model Markup Language (PMML). Producing such a model may be the goal itself or the model may be used to analyze and score additional data. As an example, a user could use Data Space to examine directly the relationship between the type of tires and likelihood of automobile accidents.

Web sites typically use Data Space services such as scoring and predictive modeling. The PSUP scoring protocol supports what is called "event driven" scoring. For example, an event, such as when a user returns to a web site, is used to update a profile about the user. The updated profile is then scored using a PMML file to produce one or more scores. These scores can then used by the web site, for example, to provide content from the web site that is relevant for the user. This is particularly important for web sites with large amounts of content or for those sites which want to personalize a user's interaction.

What we are trying to do with these protocols and industry standards is to allow you to do all the things that you need to do to create an open internet infrastructure for data. We have tried to be careful to make sure that the infrastructure scales as your data grows from gigabytes to terabytes and as your networks grow from T-1s to OC-3s. In other words, we are trying to put all the infrastructure together that is necessary to move data, analyze it, and make decisions with it in real time.

	Data Space	**Web**
Basic Objects	remote & distributed columns of data	remote (multi-media) documents
Basic Operations	a) moving data b) transforming and shaping data c) exploring and analyzing remote (columns) of data d) mining distributed (columns) of data e) scoring data	viewing remote documents
Languages	a) xml for meta-data b) Predictive Model Markup Language (PM ML) for statistical and data mining models c) Data Extraction and Transformation Markup Language (DXML) for transforming and shaping data d) ascii, XML, and SQL for data	HTML and XML
Protocols	a) DataSpace Transfer Protocol (DSTP) b) Predictive Scoring and Update (PSUP) Pro tocol	Hyprtext Transfer Protocol HTTP
Servers	a) data (DSTP) servers b) scoring (PSUP) servers	HTTP servers
Interfaces	point and click	point and click

Table 1. Data Space is modeled on the web and provides an infrastructure to move, transform, explore, and mine remote and distributed data in the same way that the web provides an infrastructure to view remote (multi-media) documents. **Table from a forth coming publication by Robert Grossman**

Diagramming the Processes and Standards

Figure One below shows this real time loop in more detail. In an off-line process, data is captured and stored in a data mining warehouse. A data mining application is then used to produce a statistical model in PMML and a set of profiles that are stored in an operational data store. In the event driven, real time environment the scoring server receives an event. Using the event, the corresponding profile is retrieved from the operational data store. The profile is updated using the event and the rules from a DXML file. The PMML file is then applied to the updated profile to produce one or more scores, which are then returned to the scoring client. Although it sounds complicated, it can all be done in a couple of milliseconds once the profile is obtained.

For example, Magnify is using this type of event driven, real time scoring for fraud detection and to personalize web sites.

Traditionally, scoring like this has been done using closed proprietary black boxes. But with the Internet, people who would ordinarily do things in proprietary languages, started to use HTML. We believe that it will be the same with the data web. Until we agreed on PMML, it used to be that all scoring, all recommendations, and all real time processes were done using closed proprietary systems and protocols. Our view is if the vendors support this, there will be a growing interest in moving completely to open standards for scoring, predictive modeling, personalization, and other emerging Data Space services.

PMML was developed by a vendor led consortium called the Data Mining Group (DMG) <http://www.dmg.org>. Currently, PMML is supported by IBM, Oracle, Microsoft, NCR, SPSS, Magnify, Angoss, Mine It, SGI, and a variety of other data mining vendors.

At the KDD conference this summer (August, 2000), the Data Mining Group sponsored a test in which several vendors including IBM, Oracle, and Magnify produced PMML files from a test set. To demonstrate interoperability, Magnify also demonstrated a scoring engine that was able to read and score PMML files from all three vendors. Since then IBM has also announced a scoring engine, and additional vendors are expected to announce scoring engines this coming year.

The DMG is now beginning to standardize DXML. The Data Space developers are also beginning the process of standardizing the specifications for PSUP and DSTP. Part of the challenge is to make sure that this suite of standards interoperate with related standards, including those for privacy (P3P) and security.

To summarize: the Data Mining Group is a very active industrial consortium focused on standards. They have been very low key from a marketing perspective and have simply focused on the standards. The Data Mining Group's members are: Angoss, IBM, Magnify, Microsoft, Mineit, NCR, Oracle, Salford Systems, SGI, SPSS, and Xchange. Reviewers: ACM SIG KDD, University of Illinois at Chicago Laboratory for Advanced Computing,

The Terabyte Challenge Testbed is also very active. In addition, a variety of industrial partners from the Data Mining Group participate in the formal demonstrations which occur twice a year. The members of the 2001 Terabyte Challenge Testbed are University of Illinois at Chicago, University of Pennsylvania, UC-Davis, Caltech, Internet 2 - Ann Arbor, NCAR, Imperial College London, Australian National University - Canberra, University of Virgina, Access Center - Arlington, and various other sites on an ad hoc basis. Basically, the Testbed creates new standards and tools for the Data Web and the Data Mining Group makes sure that it becomes vendor supported and carries on the work once the research component is minimal.

This is what I would call the twenty thousand foot over view of data space. Data Space uses open standards to provide the internet infrastructure necessary to work with scientific, engineering, business, and health care data. Unlike HTTP and HTML which are designed for multi-media documents, Data Space is somewhat more complicated because you have higher expectations when you work with data than when you work with documents.

A document you only have to read. With data you have to analyze, score and make decisions. What everyone interested in tracking and planning for the further growth and development of Internet infrastructure needs to understand is that so far the current internet barely scratches the surface of what you will be able to do with data as Data Space and similar infrastructure begins to be deployed. I'm sure that the data web will be an important driver of bandwidth over the next few years.

Some people are trying to use JDBC (Java connectivity to a relational data base) to accomplish this but it doesn't really do the job. Doing an SQL query over the web is not what this is all about. If you are an e-tailer on the Internet and you want to find out the shipping address of your customer, you would be likely to go to the Oracle database containing those addresses. The call that you make to that database could likely be a JDBC call from your Apache server.

But this is not what we are talking about. If you are a scientist and want to publish your data and let other scientists touch it — for example, are sunspots related to global warming or is El Nino related to cholera — how do you do that? Or if you are in the Fortune 500 and have data at two sites and want to personalize the web content for inbound visitors to your site, how will you accomplish this?

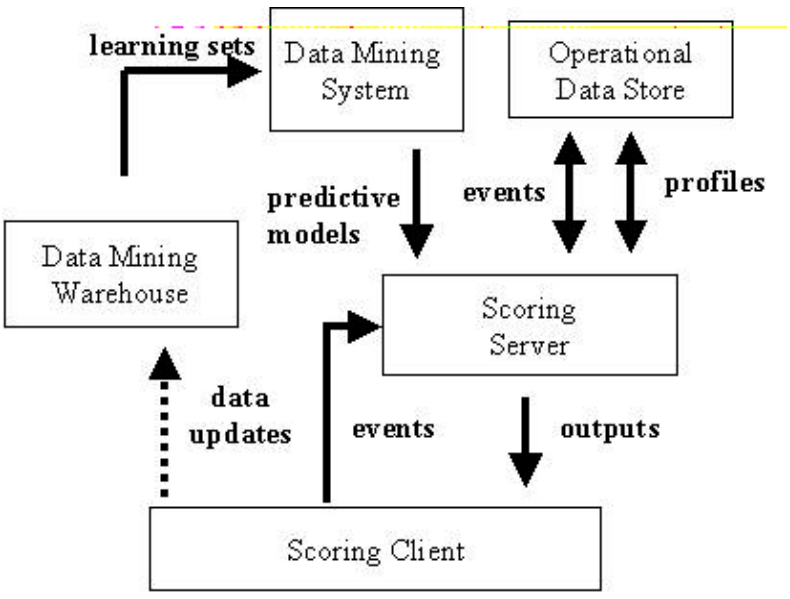

Figure 1: The Operation of the Datamining Process

COOK Report: To what extent must one worry about security? If a scientist claims to have grabbed my data, how do I know that it is really my data that he grabbed?

Grossman: That is a very real question. What we are doing is assuming that whatever protocols and infrastructure you are using to protect your documents you will also be using to protect your data. For example, if you are a corporation, you would presumably be using a virtual private network. If one wants to go even further than this, one of the master's students at the university did a quick throwaway prototype showing that digital signatures and certificates could easily be added to a data space server to provide an additional level of security and confidence.

The Next Great Web Revolution

COOK Report: What I am hearing now is that you have a set of tools built that by the time you can get them reasonably well disseminated that there will be so much that can be done that the impact will conceivably be as great as the first web revolution in spawning huge additional growth in the use of the net for new and unpremeditated tasks, and devices to enable them.

Grossman: This is our hope, and one reason that I spend at least one day a week on the phone with the Data Mining Group is that this is not going to happen without vendor buy in. It has to be done in an open way, which is why we started the Data Mining Group (DMG). We started with PMML, we are currently working on the Data Extraction and Transformation Markup Language (DXML), and we will be tackling the the Predictive Scoring and Up Date Protocol (PSUP) after that. We also worked very hard to get IBM, ORACLE, NCR, SPSS, Microsoft, and many of the smaller vendors involved. See Table 2 on this page.

The web is going to go in lots of different directions. There is always going to be the document web. Peer to peer services like Napster are here. What I am going to call Broadband Based Business Services (B3S) are coming propelled by increases in bandwidth. The result is that you will be able to move data as easily as documents. As soon as you can do this, you will want an infrastructure that will allow you to analyze the data, making decisions using the data, and to do this in real time. This is what we call the data web. E-commerce advocates talk about personalization, but today when we talk about personalization, we mean what Amazon.com does directly to books by means of a closed system. Think about what it would mean if it were open and you could create from it an infrastructure that would drive an industry. This is what we are trying to do.

There is nothing wrong with having big dreams because we know there is a lot of data out there that it basically being ignored. The amount of data is growing 6-8 orders of magnitude faster than those capable of analyzing it. This gap means that most data is being ignored. One of the roles of Data Space is to help close this gap.

COOK Report: So what are the businesses and business models springing up in this space?

Grossman: I have a simple view of the last forty or so years of commercial software. In this view, there are three eras. In the first era all software was custom. In the second era, the majority of custom software was replaced with standard applications, such as databases, accounting systems, and word processors. I see us entering a third era in which the majority of applications will be replaced by Managed Services. For example, a mail service such as Hot Mail or Yahoo Mail provides enough of the functionality of a PC-based mail application to satisfy the needs of many users.

From the user's perspective, it is much easier to use this type of managed service than to deal with the complexities of a software application. By the way, an Application Service Provider offering software on a leased base over the web provides only part of what a Managed Service Provider (MSP) provides. The key is that a MSP provides a completely outsourced solution, not simply outsourced leasing of software.

Data Space provides the perfect infrastructure for Managed Service Providers interested in working with business data. For example, Magnify is a MSP for data mining, predictive modeling, and other business intelligence services. I suspect that there will be more and more Managed Service Providers offering data analysis, data analytics, data warehousing, business intelligence, and data mining. There are already several MSPs offering storage services. Unfortunately, the underlying infrastructure for doing what needs to be done today is simply inadequate. This is where the Data Space infrastructure and related infrastructures come into play.

COOK Report: Why is the infrastructure inadequate? Not because physically you can't transport the bits? But rather it is inadequate because your software, being largely

Some of the Protocols and Languages Used in Data Space

Operation	Standard	Description
Moving Data	DSTP	The data space transfer protocol (DSTP) is the protocol used to move data between nodes in the data web.
Transforming Data	DXML	The data extraction and transformation mark up language (DXML) describes how to clean, transform, and shape data. This is usually one of the most labor intensive tasks when work ing with data.
Building Statistical Models	PMML	Statistical models are built using statistical and data mining applications. The predictive model markup language (PMML) describes the output of such systems in an open format
Scoring	PSUP	Scoring is the process of using statistical mod els to make decisions. The Predictive Scoring and Update Protocol (PSUP) is a protocol that can be used for both on line real time scoring and updates as well as scoring in an off line batch environment.

Table 2. Data Space is based upon new and emerging protocols and languages. **Table from a forth coming publication by Robert Grossman**

proprietary, doesn't transfer well from the environment of one provider to another?

Grossman: Yes. The data simply looses too much semantic information. You want protocols which retain more of the underlying semantics of the data. For example the problem is that when you move data over the web, today the data moves primarily as an HTML table, which as I pointed out earlier you cannot edit or shape in anyway. It is like a dumb paper based fax. The Data Space infrastructure would make it possible for a managed service provider to deal with data having full semantics. We want to improve the foundation on which they operate. This should accelerate their growth and accelerate the development of the functionality they can provide.

COOK Report: If I am a corporation with corporate sensitive data, am I going to have any worries about allowing a third party to apply a tool set of their own by means of which they will manipulate my data?

Grossman: Yes, but a much larger worry is that corporations stumble if they do not focus on their core competencies and outsource everything else. Very few corporations have as their core competencies data warehousing, analytics, or mining. They will therefore outsource this and when they do, they must take care to do it with an infrastructure that provides the necessary security and privacy.

COOK Report: And the companies that do this know that if they can't convince their customers that they will handle the data with proper security, they have no viable business?

Grossman: Yes. And companies have been doing this for other companies for years. All that is different is that you are replacing fedexing tapes with secure FTP (data transfer protocol) over VPNs. The security infrastructure available through virtual private networks is such that it is probably safer these days to transfer data over a VPN at an agreed upon time than it is to send it via Federal Express.

Combining Data in Science and Business Applications

But let's return for the moment to the flip side of this which is the scientific infrastructure that is pushing the data web. This is usually organized by discipline. Perhaps one of the better organized disciplines that is trying to use the current Internet to work with data is bio-informatics, where for example as they sequence the genome and they get additional information about functional genomics and related information, that information goes up on the web. In this case you can get the raw maleable data from the web but you will find that it is very hard to join that data across experimental boundaries.

In this instance, again I see some of the core ideas from Data Space as being highly applicable. Namely, how do you move data with standards? How do you transform data such as with the data transformation and extraction mark up language that we talked a little bit about? How do you score data with the PSUP protocol? What we are trying to do is to provide cross disciplinary protocols that can be layered over the existing proprietary data bases in certain domains so that you can more easily work with distributed data.

The amount of data has been exploding and what has been lacking is some underlying wide area and distributed infrastructure for working with scientific, engineering and health care data. Some of this is being addressed by the so-called "grid" style of distributed computing. For example, the NCSA (www.ncsa.uiuc.edu) and NPACI (www.sdsc.edu) computing infrastructures provide distributed computing, distributed storage, distributed security, and related services. The Data Space tools, protocols and languages that we have been talking about would by and large interoperate with grid style computing.

There are a number of start-ups which are trying to commercialize the distributed grid style of computing. *Red Herring* had a recent article about some of these companies. It is an exciting world because there is such need and such opportunity. There is a bit of a vacuum though, and, therefore I think there will be a number of protocols and languages that will compete over the next couple of years to play a role in helping grids and distributed computing to do the tasks for which they were designed.

COOK Report: Can you give some examples of what these data sets would be? How much of the likely work is cross disciplinary and do we need special tool sets in such instances? Is the cross disciplinary interaction mostly in science or is it in business as well?

Grossman: I think it is everywhere. The example of cholera and el Nino is relevant to your question. The traditional way to do something about this is to ask a graduate student to spend his or her time trying to overlay two data sets that have come from different sources and were really never meant to be overlaid. It is a very very labor intensive activity.

In the business world you might have a situation where one system records recent purchases from in bound 800 telephone orders, but when these same customers purchase on line they are treated as though they were different customers. Corporate data warehouses and some customer management systems do in fact combine information about customers across the different channels of customer interaction, but this should be really be something that is driven by the way that data is organized, labeled and moved through out an organization, say with Data Space tools.

Some of the simple ideas introduced in Data Space provide simple mechanisms that by themselves are not sufficient, but become so when combined with some of the activities going on in meta data standardization and the general move from proprietary data formats to XML data formats. Combined with these activities, the ideas of data space offer simple ways to join data from different data sources, different departments, or different disciplines in a way where their combination is formally meaningful. The units and transformations are such that they may be combined.

The commonality is generally a universal correlation key. On the customer side this would tend to be a customer ID that would tell you who the customer is. There would also probably be a temporal ID so that you would know it is the same customer at the same time. The temporal ID would tell you that on such and such day there was a web purchase while on a different day there was an offline catalogue purchase.

In a scientific domain you would say that this region of the world at this time over a period of one month had "n" many cholera cases compared to this region over the same period of time having had such and such El Nino activity. They are then formally comparable. With these examples it is possible that there could be some direct causal relationship. However there is nothing to preclude you from saying: Does this customer purchase correlate to the position of the stars at some point in time? All we can do with the underlying protocol is to provide a formal comparison using universal correlation keys that align time, place, profile ID and so on.

From this point on it is up to the analyst who you never can take out of the loop. All we are doing is providing good tools to be used by good analysts to do what they need to do in less time and less cost than would normally be the case. This is what is going on with Data Space. For example you could correlate El Nino with the stars again, because they after all have a certain position at a certain time.

COOK Report: But it would be the analyst through the use of his or her specialized knowledge who would make the assertion

and then be ready to prove that certain things could indeed be correlated in a true cause and effect pattern?

High Bandwidth and Large Data Sets

COOK Report: But parallel to this is another very interesting high performance universe. Let me turn there and ask how the work you have done with scientific data at the University of Illinois Chicago compare to the commercial development we have been talking about.

Grossman: I have been working on this infrastructure in different guises for about 8 to 10 years. Recently we have been looking at what you can do differently as you move from and a T-1 to an OC-3 connection when working with remote and distributed data. That is the DSTP protocol that I have been talking about. For small data sets this is practical today in science and business.

But to effectively move big data you need high bandwidth. Say an OC-3 that can move 10 gigabytes in 15 minutes. If you are going to examine a large data set you also need a high performance computer. You can achieve this these days by using clusters of work stations.

COOK Report: When you started to experiment with large data sets and high bandwidth, what did you learn?

Grossman: We found that if we took our standard software and ran it over an OC 3 the performance was disappointing. It soon became clear that effective use of OC 3s would not be possible without new software. There were two paths you could take. On the one hand you could spend a good deal of time tuning the network between a and b. By using jumbo packets and fine tuning the network, it is possible to move data at higher bandwidth.

But we also saw another choice — the striping of data with basic parallelism to achieve higher performance. This is the basic idea behind RAIDs (redundant arrays of inexpensive disks). For example with one commodity disk you cannot achieve very high performance. But if you used eight or 16 such disks, you could stripe the data and build high-performance disk arrays. You'd pull your first bit out of the first disk, the second bit out of the second disk, the third bit out of the third disk, and so on. We took this idea and asked what is the simplest way we could stripe data across a network connection?

Instead of opening a single socket which is how one would normally move data, we would stripe the data with the first byte in the first socket, the second byte in the second socket and so on. We soon found that we could run 16, 32, 64, and even up to 128 different sockets. This allowed us to move data from one point on the vBNS or Abilene network to another without any special tools or tuning. This allow us for the first time to saturate OC-3s with data sets talking to one another. We're now beginning to saturate OC-12s.

COOK Report: In other words rather than filling one large OC-3 or OC-12 connection, you are filling those connections with many many smaller connections?

Grossman: That's correct. We built a library to do this and called it PSockets for Parallel Sockets. We then built the basic Data Space application over PSockets. This meant that if you had a terabyte of web data in Philadelphia, and a terabyte of customer data in Chicago, you could use high-performance Data Space to move data at high bandwidth and ask basic questions about customer behavior. We did a number of experiments at the Abilene applications meeting in the spring and have run tests at the last few Super Computing Conferences using the Terabyte Challenge Test Bed.

Where several years ago we saturated OC-3s now we are beginning to saturate OC-12s and hopefully one day even OC-48s. We did this with applications and not just with moving data. A fundamental challenge today is to design and build good applications which can operate at gigabit per second speeds and beyond. We are a year or two away from this. We know that we can move date data at gigabit per second but the question is can we also build applications that function at this bandwidth? Several research groups are attacking this problem and we will build Data Space to support the efforts.

The Terabyte Challenge

COOK Report: So what has the Terabyte Challenge been all about?

Grossman: We have been running the Terabyte Challenge <http://www.ncdm.uic.edu> for the past four years and have used it both as a test bed for the basic protocols, languages and tools for Data Space, as well as a testbed for different ways to scale data intensive applications, especially remote and distributed data intensive applications. The focus has been on developing an open infrastructure for working with large and distributed data sets.

The Predictive Model Markup Language (PMML) arose in a simple way. Assume that you have data in four geographically distributed locations on the testbed and that you would like to apply data mining to each data set locally. If each data mining application produces PMML, then it becomes quite simple to combine these at one central site to get a unified view of the data. The early years of the Terabyte Challenge worked out the basic principles for PMML.

Our recent efforts have been involved in real-time event-driven decision making over distributed applications, such as PSUP and languages for data transformations (DXML). Another current focus is to build effective gigabit applications. PMML, DXML, PSUP and related protocols and languages which we are testing on the test bed form the foundation for Data Space.

COOK Report: What is the subject matter of some of the terabyte challenge network efforts all about?

Grossman: We have a Data Space node at NCAR which is the National Center for Atmospheric Research. We have another node at Cal Tech containing astronomy data. We also had a variety of financial services and business data that we sanitized and used just for purposes of demonstrating some of the types of inter-operability at our command. We also had a variety of text and data designed to understand how to improve distributed information searches on the web. Finally we had a variety of health-care and genomic data applications. The general purpose was to figure out whether the architecture standards and language of Data Space was adequate for high-volume engineering, business, healthcare and scientific use. We want to focus on real applications in each of these areas to understand better what architectures are necessary and will work productively.

We found that as we go from megabytes of data on the commodity Internet to gigabytes and terabytes on OC-3 and OC-12 networks things break and have to be re-engineered. Let me give some examples.

Mining large data means mining data that does not fit into the memory of a single workstation or even a cluster of workstations. We have built into PMML basic support for what are called ensembles of statistical models, which provide a powerful way to scale certain type of statistical algorithms to large data sets. For example, we have used PMML to build ensembles of models using 32 processors on close to a terabyte of data. I have already mentioned PSockets which supports parallel striping of data across a network connection.

In addition to using clusters for computing PMML models using data mining, clusters are also required in order to provide the parallel input-output required to feed high bandwidth applications. We are currently working on matching the parallel input-output provided by MPIO and similar software to the parallel striping providing by PSockets

so that we can build genuine end to end high bandwidth gigabit applications. This is still a fundamental challenge.

We see Data Space as a rich infrastructure that today allows for a parallel disks, parallel network striping, and parallel computation and all this is either already built or will soon be built into the basic protocols and languages of Data Space. Data Space is being designed to operate either at high bandwidth with large datasets or at lower bandwidth with smaller ones. It is designed to support both today's basic e-business applications and tomorrow's large scientific and defense applications.

The notion of a single document on a single server being served by HTTP is replaced with a cluster of workstations striping data over network connections to a remote compute cluster so that a decision can be made in real time.

This will scale down to single workstations and single sockets, but it will also scale upwards. Certainly in the grid based and the pervasive computing based methodologies that we're thinking about you must have dialable use of bandwidth and dialable use of cycles. You must be able to make decisions in real time whether they are scientific or business decisions.

Although there is a lot of work these days on quality of service, it seems to me these types of basic engineering issues are more critical just now when building gigabit applications.

COOK Report : After a follow-up interview on optical border gate way done at the Canarie meeting in Montreal two days ago I am convinced that it is very likely it will be working within a year or less. When it is and you have an entire lambda at your disposal, you can take quality of service issues and tell them to take a flying leap.

Grossman: That is possible. We are currently planning what we are calling the Terabyte Challenge Three which will be an all optical testbed. We hope that we will be able to get there in a year to 18 months. Terabyte Challenge One consisted of a terabyte of disk over the commodity network. Terabyte Challenge Two scaled up to OC-3s and today to OC-12s and OC-48s. The technical changes necessary to get from Terabyte Challenge Two to Terabyte Challenge Three will be much more fundamental than what we had to do to get Terabyte Challenge One to Terabyte Challenge Two.

COOK Report: What about applications at this new level?

Grossman: I think there will be exciting events. However, what I saw this November left me a bit disappointed in that we could not exploit OC-3 and OC-12 bandwidth more effectively. There were lots of troublesome basic issues with routers, drivers, application software, etc.

COOK Report: But if you were doing OBGP instead of paths changing all the time, you just have a few giant leaps and not many of those.

Grossman: That's true. And therefore I think that as we go to an optical Terabyte Challenge 3 we are going to see changes as fundamental as those involved in going from T1 to OC-3 bandwidth. I think that is what is going to save us. I also think that we will find that we can indeed swallow an entire lambda. We envision the Terabyte Challenge Three as an application test bed over an all optical network. We are making no secret of what we're doing and are inviting everyone to come and play.

The Processes of Data Space

COOK Report: Where are you going in your continued commercialization efforts?

Grossman: PMML is basically on track and the Data Mining Group will be releasing version 1.2 in the spring. With PMML a standard, I expect to see additional business intelligence products both exporting it and importing it. My effort will begin to shift to DXML, PSUP and DSTP and encouraging the commercialization of products and services supporting them. Data Space is just short hand to refer to the applications and services built over these and related protocols and languages. For example, although we have not talked about them standards being worked upon by others for privacy, security, meta-data, etc. will also be an important component of Data Space.

COOK Report: What has to be done then in the next 6 to 24 months? Is it primarily vendor buy in? Are most of your protocols finished?

Grossman: We have broad vendor buy in for the Predictive Model Markup Language (PMML). We are using the relevant members of the same vendor consortium that worked for PMML to get broad buy in for the Data Extraction and Transformation Markup Language (DXML). After you have the output of a business intelligence, data mining, or statistical analysis operation standardized with the Predictive Model Markup Language, both producers, such as data mining applications, and consumers, such as decision support applications can more easily interoperate. The next major stumbling bloc is transforming the data. This is currently the activity that the data mining standards group is focusing on.

COOK Report: What standards group is the DMG working with?

Grossman: Again the issue here is vendor buy in. Some times its better just to get 10 vendors in a room and have them create a de facto standard than work with an existing standards body, For example I was struck with how successful MPI is. MPI was never blessed by IEEE nor by ANSI. It was a vendor standard that changed the way people do high performance computing.

We took the same track with the PMML language and with broad acceptance by IBM, Oracle and Microsoft there was a standard. We will probably do the same thing with the Data Space Transfer Protocol. That decision will be made in the next couple of months.
COOK Report: So who will be the winner and losers? Are the only losers those who try to maintaining proprietary standards?

Grossman: I think what we will see will be similar to what happened in the emergence of the document web. There everyone focused on whether the Microsoft browser or the Netscape browser was going to win. It is where all the emotional energy went. But from a broader perspective, the interesting thing is not which browser won. Rather what changed as a result of the web was that proprietary document formats whether they be Microsoft Word or Word Perfect, or Wordstar were replaced by HTML for a broad range of applications. I think that this is what is going to happen here. There will be vendors and scientific collaborations that win and vendors and scientific collaborations that lose. I think that you will begin to see proprietary data storage, data movement and data transfer protocols replaced by open ones. With the document web whether Netscape or Microsoft won really didn't matter. What really matter was that HTTP and HTML won.

COOK Report: So single vendor, vertical proprietary kinds of systems failed and had to move over to a horizontal orientation where common protocols cut horizontally across the spectrum at every level. Is this a decent way of describing what is going on?

Grossman: It is in the following sense. When you go to the web today, the ways that you operate when you go to search or when you go to listen the MP3 or to grab real video, are all different. So you may think of them as all vertical. But they are all still driven by the underlying HTTP.

I think what is going to happen is that a managed service provider working with customer relationship management data or a pharmaceutical company working with genomic data will look quite different from one

perspective. On the other hand, both will be driven by an underlying protocol such as a DSTP or a related protocol. The underlying movement of data a the bit level and at the level of rows and columns can be handled with the same infrastructure.

Privacy Implications

COOK Report: What do you say to someone who may tend to see the purpose or at least the effect of all this technology as privacy invasive?

Grossman: That's a good question. Let's try to answer it by going back to the document web. We have of course been talking about the underlying infrastructure that will create the data web. Now if you go back to the document web, you will see that HTTP and HTML provided the foundation by which people could publish anything that they wanted on the document web. Material that was deemed offensive illegal or otherwise banned quickly appeared. Issues related to content are critical and have important societal implications; on the other hand, they are distinct from the underlying technology.

With this is background, let's look at Data Space. We have been talking about an infrastructure in which DSTP, PMML, PSUP, universal correlation keys, and all the other components we been talking about allow you to do more things with more data more easily and at lower cost. This is just going to happen one way or another, because there is a compelling need for this infrastructure.

Privacy related issues are extremely important, but distinct from the underlying infrastructure. There are no simple answers about privacy and Data Space, just as there are no simple answers about privacy and the web today. For example, some countries find some material offensive that other countries feel is appropriate for the web. There are protocols and privacy platforms being developed to manage privacy on the web — I think the best example is the Platfom for Privacy Preferences known as P3P. This provides a way for a person's browser to communicate with the server at a web site and agree upon what kind of information may be validly exchanged.

According to surveys, when someone is on the Internet, if their interaction at a particular site is improved by data which is collected at the site, three-quarters of the people surveyed are comfortable with use of data in this way, as long as the data captured is restricted to only that single site.

Compare these findings with the situations where information is used across different Web sites as with advertising or when you are in an affinity network. In these situations 80 percent of users are uncomfortable with information captured at one web site being available for use by a second web site.

I expect that this will become an increasingly important issue. A fundamental issue will be, will data collected by one web site for one purpose be used at other web sites for other purposes? I call using data from one site to improve a visitor's interactions "deep site" data and the use of data across sites "shallow cross site" data. While most consumers are comfortable with the use of deep site data, most are uncomfortable with the use of shallow cross site data.

But there is also a significant proportion of users, perhaps 25 percent, who are uncomfortable with the use of any data no matter how it has been gathered or from where. From this perspective it is responsibility of users to inform sites through P3P of what they're willing to tolerate. New browsers are including P3P capabilities within their user adjustable preference settings. While there's a responsibility on the consumer side, there is also responsibility on each of the site sides to respect properly tagged consumer privacy Information. While this is an issue in the U.S., it is even stronger issue in Europe.

Some people talk about crossing the stream as something that happens when you take data from one site and move it to another. By and large I suspect there will be fewer and fewer companies willing to "cross the stream". But I also suspect there will be more and more companies collecting single site data and exploring ways to use it to enhance the user's experience on such sites. I think the strengths of European concerns and interests in privacy and will be such that the technology infrastructure will have to comply with the general level of protection that the Europeans want. This issue will create a rather high barrier that the technology must meet.

COOK Report: while I don't mind a site like Amazon collecting information about my purchasing patterns to figure out how to serve for example a market that is under served in the way of classical CD's, I would rebel at the thought of them analyzing my purchasing patterns to see if they thought I was hooked a well enough to charge incrementally slightly higher prices on continuing visits. But I suppose this is more a matter of ethics in the use of technology then the technology per se.

Grossman: I think that is exactly correct. Evidence suggests that the majority of consumers agree that if by taking note of what they do, commerce sites can figure out how to serve them better, they will have no objections. One could suggest that the principle would be similar to going to the neighborhood clothing store where the clerk who has worked with you for 10 years knows what you want and how to satisfy you. Most people find the notion of personalization to be comfortable and most would like to be given some level of it. Where most people become uncomfortable is when that information gets to another site. I think that you will see less and less cross site activity, even though there are a number of businesses out there who are trying to make that work. On the other hand there will always be misuses and there will always be some people who are uncomfortable with any use of their personal data.

What Magnify Does

COOK Report: What can you tell me about magnify and some of its customers so that my readers can have some understanding of what the market for these kind of services is at this point? Also what is needed by way of further investment and for what particular reasons?

Grossman: Magnify is a managed services provider of personalization and predictive modeling services. One of our clients is an on-line coupon company. When someone comes to our client's site, there are probably 800 coupons that they could buy. When someone visits the site, what we do is present on the first page of the web site the 15 or so coupons that we think are most relevant to the interests of the visitor. They are free to ignore what we present. But we have seen in studies that people will click 30% to 40% more frequently when we personalize the first page because it makes the process much easier for them. They can do what they want to with fewer clicks.

It is well known that that a visitor to a web site will usually only click a few times before wandering off. A productive site will let them do what they want with fewer clicks. For this company, the 30% to 40% increase in clipping coupons immediately affects its bottom line. This is something it makes it easier to use the site because you are pushing content that the consumer wants closer to the consumer. The consumer wouldn't be on the site, if they did not want the content. In this case the data obtained is never shared with anywhere else.

COOK Report: Does this happen in the same way that when a user goes to a travel related site, the ads on that site to turn out to be travel related ads?

Grossman: Yes, but in the case you mentioned that there is a third party ad server involved. You might add that from a consumer perspective they came to that site to get web links and not ads. In the example above, we're taking the content at the site to which the consumer came to get and push-

ing it closer to the consumer. Based only on the experience of their prior visit to the site, when they get there, we immediately shuffle the content so that the content of interest is closest to the front. We are using only site specific and session specific material for the "personalization".

And to give you another example, Magnify also works with a lot of financial services companies. Consider, for example, insurance companies. They spend a lot on acquiring customers through advertising. But once they have that customer they are not very good at understanding whether the customer is getting what he or she wants. Therefore in our role as a Managed Service provider we will process all the insurance company's data on its customers. We will examine the data and, if from a statistical perspective, it looks as though that consumer is likely to "churn" — that is drop its relationship with the insurance company, we bring such information to the company's attention. The sales staff can then work to retain the customer.

Another thing we do as a managed service in the insurance field is run customer data through our systems looking for possibilities of fraud. Fraud impacts all the customers of an insurance company, not only with respect to rates, but also with respect to how long it takes to process legitimate claims. We have a scoring based method that is used to evaluate insurance claims. Most will be paid quickly but those that don't score well, will receive careful, methodical investigation. When one is red flagged, the process is handed off to a human for investigation. With this predictive modeling and scoring done and then sent over the Internet to our clients, we believe that they can process claims more quickly and accurately than would otherwise be possible. Bad claims are more likely to be found. This lowers the cost of to everyone and increases the speed with which legitimate claims will be paid.

We think this allows the insurance companies which have not been aggressive about using IT technology to become so and to leap frog ahead of their peers which have built million-dollar data warehouses and not been able to get anything useful out of them. Just because you have a data warehouse, does not mean you have the infrastructure for making good decisions using the data in it. With a managed service not only can the data warehouse be built for you, but also you can immediately begin to use the data for better decision making.

COOK Report: Are your tools such that you can pretty easily demonstrate the superior quality of this to a customer?

Grossman: Yes. For example is rather easy to show that for every dollar and insurance company spends on preventing attrition (that is a deciding what customers will leave) they will get back $27 in paid premiums.

And to sum up: I think that financial services companies and a lot of click and mortar companies are going to begin to use more and more of these services. The Data Space infrastructure lowers the cost and increases the efficiency by which corporations can deploy this kind of personalization and predicted modeling technology.

Health Care

COOK Report: if you are looking at the health care is there anything striking that stands out?

Grossman: In this area certainly one of the biggest opportunities would be to try to win by doing something interesting with the genomic sequence, protein and disease data out there. I think this is a very interesting space in which you have see for the last couple of years both proprietary genomic databases being built as well as public domain genomic data bases being built and made available. There is going to be a lot of opportunity for people to explore emerging Data Space infrastructures to develop business models that are successful and of interest to the pharmaceutical companies and to a variety of other players in that space. There is such an explosion knowledge about genomic, protein and related biological information that there is ample room for different business models.

COOK Report: If the issues of patient privacy and hospital participation could be solved, isn't one of the Holy Grails of medicine the ability to compare treatment patterns over different cross sections of patient populations in different hospitals in different parts of the country?

Grossman: Absolutely. Some treatment data has been centrally warehoused by some organizations. Hospitals are very interested in the questions of whether the clinical protocols they do are consistent with the so called benchmark clinical practices. While there is a lot of data in this area, it is very difficult to normalize the data across different hospitals because they treat populations with different social economic characteristics.

COOK Report: to get the necessary protocol's do you have to massage those which you have developed or are those products readily usable? The protocol's presumably are the basic tools and what apparently hasn't happened yet is it for enough people to become a familiar with them in begin to use them?

Grossman: Let me put what you're saying a slightly different way. The protocol's will simply lower the cost of building good analytics over clinical procedure so that you get better insight into which given procedures at which given hospitals were best for patients. There are models today, but they're expensive and hard to normalize across hospitals. People are indeed doing this in proprietary ways. All that Data Space will do is to enable them to build better models at lower cost.

I expect that the models will be in PMML because that's the most economical way to do it. Here is an example of their use. Before doing a clinical procedure, a score could be computed automatically indicating its appropriateness. If the score is too high, the doctor could be asked to evaluate again the appropriateness of the procedure. The protocol's themselves would be involved in using data from admissions, radiology, pharmacy, medical records and many other databases. Certainly a huge amount of work would be needed but we do have the basic tools and the basic infrastructure for this type of application.

Other Opportunities

COOK Report: What other investment opportunities are available to those interested in this technology?

Grossman: Today (December, 2000) the companies out there have been affected by the general dissatisfaction of the public markets with anything related to the Internet or Internet technology. Right now it's a pretty scary time out there for start-ups and companies who have just gone public. Over the coming months, as a more normal balance returns, I think there are lots of opportunities to invest in Managed Service Providers providing Data Space services, such as personalization and predictive modeling services.

COOK Report: Do you have any opinion on what the creation of these open standards will mean to investment opportunities in this industry?

Grossman: One way to look at this from the point of view of the web is that shareholder value accrued to those companies which were able to put good business models on the open infrastructure of the web. I think the game will be won by the people who build good open infrastructure for data. The good open infrastructure is really the playing field. But from an investment perspective you don't care so much about the playing field as about the team that is going to win. The team that is going to win is the one that has a good business model. What we have created with the Data Space is a good playing field.

COOK Report: But precisely what is the business model?

Grossman: There is no single good business model. One business model being tried is the managed service providers model. Another business model is to build open Data Space application software, such as for next generation web analytics or customer relationship management (CRM). Today these are being built using proprietary systems. Another business model is to build tools to support the open Data Space infrastructure itself. Finally, another business model is to supply data to the other players. I believe that there will be two or three successful companies using each of these models.

COOK Report: I gather that a lot of the infrastructure with which to build a Data Web is reasonably usable now and that over the six over the next six months to a year we're are at the take off stage of the development of a lot of interesting companies and economic activities.

Grossman: I think so.

February 2001 COOK Report

IPv6 from the Viewpoint of Mobile Wireless

Continued Cell Phone Growth to Cause Deployment of IPv6 Nets Interconnected to Non Disappearing IPv4 Infrastructure

Substantial Work Remains to Bring V6 and Data to Cell Phones

Editor's Note: Charlie **Perkins** has worked at Nokia Communication Systems Laboratory as a Research Fellow in the Wireless Internet Mobility Group since the end of 1999. He was an engineer at Sun Microsystems in1997-99 and at IBM from 1984-97. His major engineering focus has been in the area of mobile networking. He has also been heavily involved in the development of IPv6 protocols for mobile systems, and the development of network protocols for ad hoc networks.

COOK Report: From the point of view of someone recommended to me by George Abe as a middle of the road realist on IPv6, what is the IPv6 controversy all about at this point?

Part One: IPv6 Deployment Issues

Incompatibility of IPv6 and NAT Boxes

Perkins: IPv6 assumes end-to-end application addressability. The Internet was originally designed so that it would run applications with an end-to-end programming model. In other words an application at one end would interact with an application at the other end without having to deal with any third party consideration from somewhere out in the middle.

COOK Report: Something that might force the translation of the application protocol?

Perkins: Right. Or, to put it another way the designers intended that there should not have to be any discernable interaction between the application endpoints and any intermediate point. Assuming that they knew how to address one another, your application in New Jersey should be able to talk directly to my application in California. The whole point of Internet addressability has been application to application direct interaction available on a global basis. Unfortunately in the early 90s when the Internet address space became constrained, it became difficult for an application to find out the IP address of the counterpart to which it wanted to link because that counterpart might not even have an IP address that was globally visible. This came about when enterprises began to insert firewalls and packet filters surrounding the networks within their domain of administration. The constrained address space also caused the reorganization of the address space into CIDR in order to make more appropriately sized subnets available as the Internet began its rapid expansion. From the point of view of the reduction of address consumption in IPv4, CIDR has been dramatically successful.

The unfortunate side effect of all this is that people have been looking at other ways to obtain addressability. The largest successful effort in this direction has been in the area of Network Address Translation boxes because now people can install a NAT box at their network boundary and get a lot of addresses back. Unfortunately those addresses are more or less localized to the domain behind the NAT box. You can get outside only by running something that the NAT box knows how to do translation for at the application levels.

In such a case it will give you something that appears to be an IP address from the outside world. If it is an application of which the NAT box is aware, this will happen automatically. The other constraint is that the NAT box can do this only with outbound data. An inbound application query goes unrecognized. You cannot have a server or peer elsewhere in the Internet suddenly strike up a transaction or conversation with any nodes that lie on the inward side of the NAT box. The node behind the domain has to initiate anything that takes place.

COOK Report: In other words if I had a publicly advertised static IP address, the node behind your NAT box could initiate a transaction with me. However, if I too lay buried behind a NAT box on my end, and conversation between us would be out of luck?

Perkins: Yes. I can, however, imagine special solutions where I may say I'll call you on the phone to set up some means of rendezvous between our machines. The NAT box serves as an application level gateway. You could have a remote box doing the same sort of thing and setting up some kind of encapsulated transmission. My point is that with some effort solutions could be obtained. The problem is that such solutions would not be easy and that they are also not scalable.

COOK Report: Aside from these problems, how do you cope with the criticism of people like Sean Doran who say that there are so many NAT boxes standing guard over untold numbers of corporate networks that the cost of replacing them with networks reengineered to do IPv6 would be prohibitive when compared to negligible benefits available from IPv6? Have you ever seen any one who has attempted to figure out how ubiquitous the NATs are?

IPv4 and IPv6 to Co-exist

Perkins: They are very widely deployed. There is no doubt about that. But on the other hand I think this argument is not relevant and I will tell you why. We are going to have a lot of IPv4 nodes out there. And we already have a lot of and NAT boxes out there. And this is just in the way it is. And these conditions are not going to change. But what will change is that we will have to have more and more IP addresses. So those nodes using IPv4 will continue to do so. However as you look at the question of IP address ability for more than one billion cell phones that are likely to be deployed over the next 18 months, you begin to face the question as to whether IPv6 is a possible solution if you want to have this end-to-end network connectivity. Now if you can sell it long and hard and manage to convince enough people that they really don't need that capability, you may be OK. The only problem is that the phone companies are not really trying to do this. They are trying to make obviously a very very accessible system. The problem then becomes the fact that NAT stands in the way of this objective.

COOK Report: So as someone who stands on the cell phone side of things, what are the changes that need to be made to get from where we are to a network of a billion or more IPv6 accessible cell phones?

Perkins: You know it's always impossible to predict the future but I can make some reasonable suggestions about it. So first of all, at a fairly abstract level, this business about insufficient addressability in Ipv4 pretty much causes you to go to a NAT style solution. This applies even when you are

slicing up the addressing domain in IPv4 and IPv6. Therefore in the future it seems likely that we will have a lot of IPv4 nodes, and very likely even more IPv6 nodes - a number that will grow steadily over time.

COOK Report: But the network administrators are not going to throw out their NAT boxes any time soon?

Perkins: This is quite true. They will stay until more people involved with the maintenance of each node decide that it's more of a hassle to keep it than to change it.

COOK Report: Let's look then at large corporations like and Hewlett-Packard are Sun Microsystems or Nokia. These companies have and very large IPv4 numbered networks. When the "hassle factor" arises, are you saying they will have to begin to renumber?

Perkins: No. IPv4 and IPv6 can co-exist in the same general network because they do not collide with each other. They just have to know how to address each other. For example you can have a router that routes IPv6 packets and IPv4 packets on the same network. If you want to have some IPv6 addressability, all you have to do is get some of these routers. The packets will start flowing and everything works exactly as you would like for it to work, until you get to this one question about how do I get one of my existing IPv4 boxes to communicate with one of my new IPv6 boxes?

COOK Report: I want to hear the answer to that. But first about routers that route both protocols. Would you expand some more on who makes them and how they work.

Perkins: In Cisco I understand that IPv6 is obtainable mainly through a software upgrade and that in the future they will be doing some degree of hardware support as well. As far as Nokia is concerned we have this IPRG platform. While it is not a router, it does fire walling with check Point Software but also handles IPv6. As far as Nortel is concerned, I don't know the details of their product line, but I do think they have active support for IPv6.

COOK Report: So what you are saying and that the Nokia I P R G platform is a device that may be used to substitute for an existing IPv4 firewall and NAT box combination with the additional capability that it will allow IPv6 to go wherever it needs to go?

Perkins: Yes. While I am not an expert on this particular product would you suggest is certainly generally correct. You may also use IPv4 routers even when you have no IPv6 routers because then you have tunneling solutions that are available. There are about seven or eight different techniques that you can use to handle IPv6 in situations like this. These are in discussed in the IETF and a working group known as the NGTrans group. Go to http://www.ietf.org/html.charters/ngtrans-charter.html Once on the page you will see the following listing of internet -drafts:

The 6bone registry (21895 bytes)
Overview of Transition Techniques for IPv6-only to Talk to IPv4-only Communication (17839 bytes)
Connection of IPv6 Domains via IPv4 Clouds (55716 bytes)
IPv6 Tunnel Broker (24546 bytes)
An overview of the Introduction of IPv6 in the Internet (72214 bytes)
A SOCKS-based IPv6/IPv4 Gateway Mechanism (22633 bytes)
Dual Stack Transition Mechanism (DSTM) (37829 bytes)
IPv6 over IPv4 tunnels for home to Internet access (9138 bytes)

Also you will find relevant material in the RFC list.

Routing Aspects of IPv6 Transition (RFC 2185) (31281 bytes)
Stateless IP/ICMP Translation Algorithm (SIIT) (RFC 2765) (60924 bytes)
Network Address Translation - Protocol Translation (NAT-PT] (RFC 2766) (51015 bytes)
Dual Stack Hosts using the Bump-In-the-Stack Technique (BIS) (RFC 2767) (27133 bytes)
6Bone Backbone Routing Guidelines (RFC 2772) (29352 bytes)
Transition Mechanisms for IPv6 Hosts and Routers (RFC 2893) (64358 bytes)
6BONE pTLA and pNLA Formats (pTLA) (RFC 2921) (14218 bytes)

COOK Report: What you are saying then is that for someone who wants to get IPv6 up and running, one of the best things that can be done is to go to that web page and read some of the solutions described in the URL s above.

Perkins: Right.

Communication Between IPv4 and IPv6

COOK Report: This leaves us with a situation where both protocols can be currently used. But if one is to do this, one has to be able to have IPv4 be able to talk to IPv6 and vice versa.

Perkins: That correct. So let's say that while you have a growing population of IPv6 nodes, you may also expect your population of IPv4 nodes to grow as well until they start running out of addresses. And this will lead to a situation where you have an IPv4 section of the Internet which is very very large but losing in overall market share.

As I was going to say before, this NAT solution is a very natural solution when you are in a situation where you have to slice up your address space. And if you think about it in this fashion, having large domains of both IPv4 and IPv6 is merely one way to partition the possibility of overall IP address space in general. In such a situation with the right kind of fire wall NAT platform, you can even mention translating IPv6 addresses into IPv4 addresses at the border of the domains so that IPv4 applications can in effect be tricked into believing that what is going on is only an interaction between two IPv4 applications.

COOK Report: And this is because it is going through an address translation barrier?

Perkins: Exactly. Except in this case you are translating from 32 to 128 bits instead of from 32 to 32 bits.

COOK Report: And yet there are more than a few purists out there who, when faced with this explanation, would likely say that it is inadequate. That nothing will do until all IPv4 applications are swept from the path and the network has IPv6 end-to-end purity once again. Would it be safe to assume that this is just absolutely not happen until the administrators of your various IPv4 parts of the Internet decided they have no choice except to close down? I ask this because I think there's an implication from which you just said that you may have some v6 applications behind fire walls and still not be able to do end-to-end? Am I hearing you correctly?

Perkins: No. Because the IPv6 routing could just occur naturally and when it has to be translated into an IPv4 address, it goes through the firewall which does packet filtering in exactly the same way whether or not the packet is an IPv4 or an IPv6 packet.

COOK Report: In other words IPv6 normally functions as an end to end service unless it has to be translated in order to communicate with an IPv4 address.

Perkins: That is correct. Another way to look at it is that the entire IPv4 address space of a quadrillion numbers is just a tiny speck in the over all vastly larger IPv6 number space. You can take any IPv4 address, put zeros in front of it and wind up with an IPv6 address. Your routers then can route those kinds of addresses. If any application in v6 land wants to interact with an address that looks like a lot of zeros followed by 32 bits of v4, the routers can handle that by routing them directly toward the appropriate NAT box. This kind of solution becomes available anytime that you can set the routing

tables up. Unfortunately, it doesn't really become an important solution until the total size of the IPv6 routing tables has become comparable to the size of the IPv4 tables.

COOK Report: How would you describe what you see coming in the next year or so in the sense of a growing number of IPv6 nodes and the application within the network of more and more devices that can handle both protocols at once?

Perkins: It will continue to spread but it will take more than a year. The whole thing about IPv6 to begin with was to develop a protocol, deploy it, and do what IETF does well which is to get to inter-operability testing going and then to just start to build it. People want to buy solutions to the problems facing them, be they IPv4 problems or IPv6 problems. People will want to buy solutions for their IPv4 problems and for their IPv6 problems. Eventually the solutions for the IPv4 problems may become more expensive than the solutions for the IPv6 problems. This will be true in part because the IPv6 solutions that are already available will become cheaper as IPv6 grows in market share.

COOK Report: So where does the imminent arrival of a billion IP addressable cell phones fit into this picture?

Data and Addressing in the Cell Phone Inundation

Perkins: The answer is that there will be a billion cell phones within the next 12 to 18 months but that by no means will all of them be IPv6 addressable. Increasingly your cell-phone-using population will expect their phones to be data addressable. Now the number of times today that I expect to get data on my cell phone is usually about zero. This is because when someone uses the phone they do it for purposes of circuit switched voice. But, as time goes on, the data will begin to be more and more important. Moreover voice will increasingly arrive as data — in other words as a voice-over-IP. People will begin to understand new ways in which these devices can be even more useful if they have Internet visibility. To do this they will need Internet addressability in order to be able to poll the applications that people are used to using today. The question is not whether this will happen but how fast? There are a lot of predictions for the very rapid emergence of an increase in market share of voice-over IP. Although I'm not a marketing specialist, from my perspective, it seems like it this may take three to five years before voice over IP is able to capture a 50% market share with a circuit switched voice.

COOK Report: What is slowing it down?

Perkins: For one thing, today, when you make a telephone call, you expect it to work. The quality of service issue is also a big drag on generalized VOIP usage. Having to go across network boundaries means that voice quality cannot be guaranteed. According to my understanding, the best market for VOIP now occurs when an ISP has extra capacity. In such a case it may then take large chunks of circuit switched voice and route such traffic over its IP networks.

COOK Report: To make such a rollout occur, one of the things that must be done is to hand out IPv6 address space is it not?

Perkins: Addressing is certainly a big deal. Various set-up mechanisms must be elaborated to make the addresses available. Over time people will begin to see the benefits and to request it more. In Japan this I-mode capability is pretty successful and there are a lot of rosy predictions for WAP as well.

COOK Report: But from the point of view of Nokia, what are you guys going to have to do to get IPv6 as an integral part of your cell phone capability? Don't you need to get some blocks of IPv6 address space in order to put individual IPv6 addresses into your phones? I believe the routing Registries have policies for assigning IPv6 space. But isn't one problem the fact that they're planning on charging for it in the IPv6 address blocks?

Perkins: I really don't now any of the details of this particular set of policies. But I do know that when very large companies step forward and ask for a bloc of a billion IPv6 addresses, if the registry can't make such an assignment available on a reasonable basis, very large companies will find a registry that can.

COOK Report: So is undoubtedly has something to do with the stories that I hear about a cell phone industry demanding a reasonable allocation policies on IPv6 as well as follow-through within the next six months or so? Given the amount of addresses desired and so-far-announced policies on pricing might not the result be a cost impact rivaling the amount of money spent on the spectrum auctions by the industry and isn't it clear that such an expenditure would be unsustainable by everyone involved?

Perkins: Spectrum auctions are governmental regulation of a limited resource. Given the almost unlimited number of IPv6 addresses, I don't see any way that the government could attempt anything similar. If you send me a packet that has some numbers inside, and I route it to the next place online, that it is just a logical thing to do that isn't easily amenable to government regulation.

And there was a discussion a few months ago regarding the size of the IPv6 prefixes that ICANN should handout. What resulted was a number on the order of a /48 as a prefix for an IPv6 address block. This is a huge number with plenty of addresses. Anyone who wants to start charging for its use will run into trouble. If anyone on the downstream side of ICANN tries to impose a charge, it seems to me that the market competition would render such a model unsustainable.

COOK Report: If you assume that the era of large-scale use of IPv6 in addressable nodes for data transmission is 24 to 36 months away, one of the things that has to get done to get from here to there is figuring out how to handle some of the address registries allocation policies and we just talked about. Now let's take it for granted that this does get done satisfactorily. What else needs to happen? For example where does the support of Microsoft operating system platforms fit in as a prerequisite?

Perkins: You can go to their web site and download right now an IPv6 stack for your computer. When they're going to sell an entire operating system that runs IPv6 is an unknown question.

Part Two IPv6 and Cell Phones

COOK Report: As far as IPv6 is concerned then what are the areas then that occupy your time and concerns?

Perkins: I spend all my time in the engineering area working on mobile networking stuff. Right now IPv6 is being standardized into various telephone protocols. In Europe you have 3GPP. This stands for the Third Generation Partnership Project. It is a standards group that has been working in this area for some years. They have said you must have IPv6 addresses in your 3GPP telephones for, I believe, release 00. And this release is basically complete. We next have to start building it and when we do this, you shouldl expect to see things that run according to the standard within the next 18 months to two years. By building it is meant both putting it into the cell phones themselves and into the cell phone supporting infrastructure.

The assertion that a generation of cell phones is running IPv6 is a subject to a large degree of interpretation. It says nothing about the performance characteristics of the product. Consequently, right now we're looking at ways to achieve, in a single product, mobility combined with sufficiently good performance in a way that people can feel that they have the quality of circuit switched voice in a product founded on IP mobility. Note that

doing voice-over IP in a wire line network is quite different from doing it in a cellular network. The constraints are much more difficult than those encountered in a wire line network.

COOK Report: I see the concerns about moving from wire line to cellular in the use of IPv6 but tell me whether there is an issue of building a cellular data distribution capability into a network with a choice of either IPv4 or IPv6? Is the situation that in this case you assume IPv6 by default?

Perkins: It's difficult for me to answer a question like that because we're working only on IPv6.

COOK Report: Can it be assumed in that any decision to go IPv6 only is because of scaling issues involving vast numbers of users?

Perkins: We are going have a lot of telephones involving also voice and data. Therefore if you want to do those two things together the use of IPv6 becomes the only reasonable alternative.

Using IPv6 to Bring Data to Cell Phones

COOK Report: What then do you have to do to achieve your objectives?

Perkins: Let's put in this way. If you could say that there are perhaps a hundred problems, then you can say that we've looked at a half-dozen to ten of the total. We have answers for most of what we have looked at but, as we look, we see more and more problems. For example there are a lot of problems in security and a lot of problems in quality of service. There are also a lot of problems in header compression. Also the way in which the base stations are coordinated to manage spectrum most effectively is historically not very friendly towards the IP model.

All of these things add up to a situation where, as I mentioned before, you can employ IPv6 now. But for specific applications like voice IPv6 cannot currently match the performance of analog voice-over the air as a part of the PSTN. Now we're going to change this and believe that we will be able to equal or exceed the current capabilities of analog voice over the air as part of the PSTN.

COOK Report: So would you be willing to outline more methodically where you will have to go in order to achieve that objective?

Perkins: The first thing must be done is the management of header compression. If voice packets are only about 25 bytes with IPv6 headers they become 85 bytes. You must really do something about the size of those headers, for if you do not you're throwing away nearly two-thirds of your bandwidth. Fortunately IPv6 has very compressible headers and it is possible, depending on the exact situation to get to the size of your overhead down to between one and three bytes per packet.

COOK Report: Is there a specific working group doing this?

Perkins: Yes. It is the ROHC (Robust Header Compression) Group. The reason that it must be robust is that you are dealing with a lossy interface. If you operate in an environment where you don't lose packets, header compression schemes can be much easier. Header compression works by throwing away as many fields in the header as you can. The way that you can throw away a field is by keeping track of what the current state of the transmission is. For example if you want to avoid having to send the IP source and destination addresses and you are operating on a known radio channel, and then you can compress a great deal by not having to send the IPv6 addresses with each packet.

You have some more header fields of various kinds that you have to deal with. But some of these are the same, or continuously varying in every packet. Of those that change with time some change slowly, some rapidly and some are unpredictable. For IPv6 most fields are highly predictable. For IPv4, the fields are slightly less predictable.

You may have a field that increases by one with each successive packet. You can pretty much get this field perfect by not sending it and incrementing its place holder by one with every successive packet on the receiving end. However, if you unexpectedly lose a packet, your algorithm is hosed. Therefore you must take particular care to make your algorithms robust.

COOK Report: Could you fix this particular problem by using a check sum?

Perkins: There is more than one technique that would work, but yes there is a check sum that you could use and already there are some check sums at other layers in the protocol. As a result this is manageable.

COOK Report: So one of your concerns would the in figuring out how to handle the headers under different kinds of operating conditions?

Perkins: You make a very interesting point there. Because if you could just establish a transmitter and receiver connected by wire, you could eliminate all problems stemming from packet loss. If you take the wire away you can still do a good job of header compression but you do have to worry about losing packets. However, if you add mobility to the removal of the wire, things become really different and you have to begin doing things like adding coordination to your points of attachment. Nevertheless, we have solved the problem. We've built it. We've shown it. And now we're involved in standardization.

COOK Report: In other words right now a Nokia phone could use what you developed to talk to another Nokia phone but without the standard there is no guarantee that a phone made by someone else can talk to the Nokia product and vice-versa?

Perkins: That's almost correct. The real compatibility requirements are between the telephone and base stations. By the time you have done the first hop you have already made the compression and decompression. When the package shows up on the network side, it's decompressed and is carried as a part of the normal traffic supported by the global public network infrastructure.

COOK Report: You are talking then about the example of someone who might want to phone from a car in New York City to Chicago or San Francisco or Hong Kong and, on arrival at the destination city, be able to hop out of the PSTN infrastructure and be delivered locally as a cellular call?

Perkins: That's it.

IPv6 and Packetized Voice

COOK Report: What then would be the purpose of using the IPv6 protocol in a situation where you were doing primarily cellular within a metropolitan area network or between such networks if your call was not transmitted over the Internet to be completed?

Perkins: What is good about IP is that it allows you to transmit packetized voice. People doing packetized voice care more about the packets than about IP. But eventually people are going to want to move packets and voice-over IP packets onto their networks.

COOK Report: When you say "eventually," if it were like five years from now, must we assume that all of digital telecommunications will be IP?

Perkins: Yes indeed. A lot of people assume that telecommunications in general will mean communication by means of the IP protocol and that we will get to this state in rather less than five years.

One thing to recall about current voice technologies including wireless is that everything you do is geared to knowing how much bandwidth you need in order to make the voice that you want happen. Consequently when you make a phone call over the air, a lot of effort goes into giving you a time slot or code necessary to give you a radio resource. Now if you want to do some data in addition to voice or if you expect burstiness in your transmission, the old algorithm keyed to the reliable delivery of voice only goes very quickly to hell in a hand basket.

If you want to make any convergence between voice and data, it makes sense to use IP to accomplish the task. The current PSTN this is not geared to do that sort of thing. For proof of this statement you may want to take a look at the current "success" of CDPD. Being able to mix the use of voice and data together necessitates the use of IP.

COOK Report: So what you must be prepared to do is to open up an IP channel at some point in the future between two people where the people and devices could be using the bandwidth available to them by separating it into a voice channel and the data channel on a moment to moment basis.

Perkins: That's correct. This is the statistical multiplexing of IP does so well.

COOK Report: OK. But if we're talking about how we get from here to there after header compression, what would be next?

Perkins: We have implemented a buffering scheme in order to handle a situation where you are momentarily disconnected as you move between points of connection and are handed off from one base station to another. When you are handed off from one base station to another, it is very easy to lose packets. Packet loss for voice handoff is one situation. But if we ever do go to handing off video, the situation becomes even more demanding. Whatever the case, in order to have really seamless mobility we have had to commit ourselves to buffers. We have done this and it is working pretty well.

COOK Report: Is this something that is being done in a working group or does it not need to be done there?

Perkins: We have submitted an Internet draft for consideration. Now if you move to a new place in the network and your buffers are somewhere else in the network, then you may be responsible for requesting delivery of those buffers to the point where you have shown up. In our particular kind of a system you will have protocol requirements. In other systems, the network does everything for you. It can help figure out where you're going, where you have been, and how many of buffers you have traversed in the process. If you have a situation where the intelligence is at the edges of the network which is there primarily just to help push bits around, you will need a protocol so that you can tell the network where the bits are buffered.

COOK Report: And at such a point in the capacity for doing the buffering would presumably be in a device on your end and in another device on my end?

Perkins: Actually not because you would then have an Internet separating us and we want to have these buffers delivering their content to each other very,very fast. Now if you set the buffering at the end points you might just as well call it TCP, because that's what TCP does right?

COOK Report: OK. But another model for describing this that comes to mind would be the client-server model where the server would be the device containing the buffer and it becomes up to the server to get it the contents of the buffer out to the clients at the appropriate times. Is that a reasonable analogy?

Perkins: It is sort of reasonable at a particular level of abstraction if you see the buffering that your telephone is relying upon as the "client" for the phone. But if you look at the other side of what it means to be a server (as something to log into) you cannot afford to take the time to do this and the analogy fails. The point is to save as much time as possible and the way to do that securely is to have some protocol and communication features involved.

COOK Report: OK. Duly noted. Now where does IP SEC come into all this? All IPv6 implementations are required to include IPSec in order be considered compliant. Are they not?

Unsolved Security Issues Yield non Standard Kludged Systems

Perkins: Yes. If your wireless device shows up somewhere on the network and makes a request to do header compression and relocation or to receive the contents of its buffer or asks for its quality of service parameters to be transferred to its new position in the network, these actions must be taken securely by means of authentication. Otherwise people can take advantage of and abuse the system.

COOK Report: In other words aren't you turning against the possibility of somebody being able to listen to the bit stream and scoop it out of the air putting something back in its place that is different than what the parties expected? You are guarding against my receiving something different than what you sent me.

Perkins: Not quite. What you're talking about is different. Let's make the distinction between privacy and authorization, because what you're talking about is a privacy issue and what I was talking about is an authorization issue. I don't want anyone to be able to steal my identity to the point of having packets delivered to the wrong place. But you could imagine packets going exactly to the right place with no privacy at all because any one could overhear what was being sent and received. The issue about making secure handovers is an authentication and authorization issue.

COOK Report: And what we're talking about here is a by product of the issue of mobility where you have to hand over the signal from one base station to the other? Now can IPsec be used to accomplish the necessary protections and, if not, what does have to be done?

Perkins: That is actually a pretty complicated question and there are a lot of moving parts to it. For example just a situation that you're talking about with a couple of base stations, a couple of mobile nodes, a packet stream coming in from another place, and the routers upstream from the base station trying to figure out where to send the packets. But in the way you have laid out the issues of deployment there are other entities farther into the Internet and into the core networks that start having to have a role to play.

All these parts of a mobile IPv6 Internet have particular kinds of security and authentication needs and the way that this is to be built up is, I think, the biggest unsolved problem. All that I said before was that you could take some of v6 boxes and put them together in order to build a network in that would indeed work.

But now let's take a situation where you are a service provider and ask whether you, as a service provider, can build up a v6 network and expect it to work — in other words be able to sell service to customers. And the truth is that you could only do this in a fairly non standard way right now because the protocols for doing all the authorization and mobility management in IPv6 are simply not finished. You cannot buy equipment that will do this. I would also guess that it would be pretty good number of months before you could begin to do in any way shape or form what we're talking about. And when you are talking about this kind of engineering you must ask: is it being done for normal data, or for normal data and voice?

COOK Report: And from convergence point of view you want to be able to do it with both voice and data — a goal which presents us with the most difficult technical hurdles.

Perkins: It is the most difficult because voice has real time delay requirements. Just to do the routing in the context of real-time constraints is a difficult challenge. But when you look at doing this and also look in doing it securely, you find that the problem begins to partition itself into easy parts and into hard parts.

COOK Report: Do you see it being done first without regard for security or do you see it being done only in a if security is also handled as an integral part of the original process?

Perkins: The problem is one of managing the distribution of the keys. To hand out all the keys to all the base stations and to all local nodes and other entities there that have to interact is not an easy task. It can be done and I can draw pictures on the board of how to do it. In fact we have already done quite a bit of that, but to do it with appropriate performance and with standards is very definitely non trivial.

Immediate Architectural and Security Issues

COOK Report: What is it both possible to do in the shorter term without doing that and possible to do that from an economic point of view is worthwhile doing?

Perkins: You can put up a network of IPv6 base stations and applications. IPv6 already has encryption headers and authorization headers, all of which work well in a static wireless network. But with a cellular infrastructure you have a mobile computer coming into an area where it doesn't have any pre-existing security associations. It will have to establish this by some sort of protocol with a remote agent. Providing static security associations within an enterprise network is a problem of another color entirely.

COOK Report: So, in other words, the difference is if I have a cell phone and I am out in the ballpark somewhere, where I am staying for a couple of hours versus wanting to transmit my data from a car moving at high speed down the highway. Is that part of what you're talking?

Perkins: Yes. It's a kind of middle ground because there is not as much mobility involved.

COOK Report: Or are we thinking of a situation where there is a telephone network and I am using one device plugged into the PSTN? Moreover this device is a non mobile device with only one number but I am going by means of IPv6 over the Internet rather than by means of ordinary circuit switched voice over the PSTN?

Perkins: All these things are possible and I'm sure they are all going to be done. But let me also see if I can answer your original questions. The point at which you show up at the ballpark is the point at which you will have to do authorization. Authorization presumably will be done by interacting with the network as administered at the ballpark. And once you get authorized you are pretty well established for your use of the network while at the ballpark. Nevertheless you cannot ignore the existence of a discontinuity in your use the network. In other word something had to change from your prior use to your use while at the ballpark. The point is you should be able have this change in authorization happen automatically. But once you get there and it's been done, there wouldn't be any further issues of mobility management or re-authorizations.

COOK Report: In words all your devices must be geographically aware of where they are and be able to revise that awareness on the fly?

Perkins: Yes. Except I would not use the word "geographically." Everything depends on network connectivity which does not map very well to geography.

COOK Report: But I was thinking geographically in relationship to their position vis-a-vis the nearest base station in the network with which they must check in.

Perkins: Exactly the base station will be the point of contact with the Internet and the point at which the protocol begins.

COOK Report: One should think of a two-part communication process? Communication between the person using a mobile device and the base station is one set of processes and communication between the base station and the rest in the Internet is the other? You could consider this the middle if you imagine that the call will eventually reach another base station and pop out on the other side before going to its recipient.

Perkins: Yes. And what is in the middle is very likely group of service providers that expect to get some money for the transaction.

COOK Report: Are we looking at the same situation here that we looked at in all the classical quality of service issues namely one of crossing provider boundaries. In other words if I'm a UUNET and UUNET is so big that I can get almost everywhere else in the world I have an easier time of things tha if I have to cross over from UUNET to a recipient for my call that happens to be a Sprint customer. Is that correct?

Perkins: Yes. That is right.

COOK Report: But to get the commercial viability which you want, you at Nokia have to be working people at Ericsson, Motorola and others on the technology and economics of crossing provider boundaries. True?

Perkins: The operators are clamoring for the ability so that usually produces good results from the equipment manufacturers.

COOK Report: Can you explain what has to be done?

Authentication, Authorization and Accounting

Perkins: Yes. I am pretty well involved in this at Nokia and within the IETF. Right now, there is an Internet protocol called Radius. It is used by dial up users to authenticate their connection to the network. As a result the process of doing this on a dial up and login basis is pretty well understood. Right now within the IETF there is a working group called AAA (authentication, authorization and accounting).

Since Radius only works for static objects and has some other difficulties as well, this AAA working group is building up a replacement protocol for Radius. As a result, we will have this AAA protocol that will come with features such as session measurement and Accounting. Tied in with IP SEC, AAA will do authorization and accounting for services such as mobile IP. And if you say that mobile IP is just a part of IP, then what you end up charging for is IP.

You will have then these mobile devices that need to interact with the AAA agents. They will do this according to AAA protocols. They will be simple protocols but ones that have to be done very securely. All this means that there must be a connection between the AAA working group and the mobile computing working group in order that what gets done does so in a way that is interoperable, modular and scalable. Our current challenge is focused on the workings of these two groups where most of the issues that arise have to do with key distribution.

Let me give you an example. If you move from your own network to another networks and you receive a key, that key must come from your home network because your security association lies only with your home network. So the key that is sent from your home network to the visiting network must

be different than your actual key because you don't want anyone else to see your actual key. Therefore it must be some sort of derived key.

COOK Report: Is it like a public-private key?

Perkins: It really is a symmetric key. But you could also map it as public-private and in such a case the important part of course would be the private key. In general you want the appropriate key to follow you from your home network and show up at the proper spot in the visited network. But if it shows up at where you are, and then you move to a new place, you want that key to follow you and become available to you in that new position without having to go back to the home network all the time in order to retrieve it. What you have to do boils down to a combination of encapsulation and encoding because, when you're doing these kinds of transmissions, it is no longer safe to assume that everyone is trustworthy. All transmissions must be done in such a way that privacy is insured.

This then is a thumbnail sketch of what needs to be done and is being done. A good part of the story is that this process is basically complete for IP version four. What we are in the process of doing now is mapping this into an IPv6 solution.

COOK Report: So I'm hearing from an analog, circuit switched, cellular point of view that everything is quite mature and in that everything works. And but I am also hearing now that if you want to converge voice and data in and use IP to do that, things are rather less far along. From an IPv4 point of view we have things working with that cannot yet be done in v6 perhaps in part because not many people are doing v6 yet? Do people have to do everything first in v4 and then transfer the result to v6? Or can they begin to do things in directly in v6?

Perkins: People are beginning to do things directly in v6. And, for the most part, development is proceeding in parallel in both v4 and v6.

COOK Report: So there is then a very decent market for the use of IPv4 in mobile communications?

Perkins: Certainly. There are a lot of people who want to use v4. But I don't think we will ever get to global deployment of mobile IPv4 for voice-over IP. I think by the time voice-over-IP really comes into play, we will be using largely IPv6.

COOK Report: Then IPv4 mobile use is primarily data rather than voice-over IP?

Perkins: I think that is generally correct. But on the other hand, I don't have any set of specific data to back up that claim.

Summing Up

COOK Report: How would you begin to sum up where we are?

Perkins: We've covered most everything except for quality of service. Quality of service is an issue and it has some difficulties which, unfortunately, I'm not all expert in. But with that caveat what moving from one place to the other in the development effort for wireless cellular IP seems to boil down to is that it appears that there is quite a lot of capacity within the core network infrastructure. What this means is that when you move from one place to the other it that you don't cause difficulties with the current allocation of spectrum. The thing that you must avoid is that when you are making a telephone call and move to a new place in the network where there are already 10 people making telephone calls, you don't want your new telephone call to crash theirs. Therefore it is a lot better to give you a busy signal than to cause their voice sessions to degrade. The protocol necessary to handle this could be called Network Admission Control. Local admission control algorithms are needed and then these admission control algorithms can be aggregated into some sort of a back hauled differentiated services solution.

COOK Report: Would you try to sum up for my readers in a few sentences the difference between the world where we are now with voice and mobile devices that may not use IPv4 let alone IPv6 and the world where we are headed with all the new technical development underway?

Perkins: First of all the world is going mobile; that's crucial. So what ever is done protocol-wise must work with mobile as well as with static devices. Second of all the world is going Internet. However, we have a long way to go there because right now the vast majority of voice traffic is not yet on the Internet. Also there is not yet a good integration between data and voice. But eventually there will be. The third piece is multiplying by ten the number of Internet devices in use over the next four years or so. This means that we are going to have scalability problems and more and more customer demands.

Right now the majority of Internet deployed devices are not mobile but the time is coming when the majority will be mobile. This is in part because of the nature of the devices themselves and because practically all the new devices coming on line will be mobile. Remember that each of us, as we travel, will have an with us the power to communicate with not only home and office computers might also with a myriad of static condition monitoring devices in both home and office as well as within environments through which we pass. All these factors taken together over the long haul make an almost unbeatable case for IPv6 and for mobile IPv6.

My guess is that in terms of standards we are 80 to 85 percent of where we need to be and that in terms of equipment availability we can buy most of the pieces that we need today but not quite all. And that it will probably be three years before IPv6 is viewed as having a significant market share. I think that it will be another four to five years before the market share of IPv6 is more than 50%.

COOK Report: Are you saying then that the criticism that switching to IPv6 will be too costly to do is the wrong way to look at it? But that to the contrary there are reasons to begin doing IPv6 that will not cost a lot money and will not entail huge expenses of translating from one system to another? And that the reason to start down an incremental adoption path is that it will just grow and grow and grow independent of anything else that is happening?

Perkins: I think that's fair point of view. To pose the problem as one of how we're going to change all the existing devices is just not right because people don't want to change existing devices and a lot of benefits of IPv6 can be gotten without changing existing IPv4 devices.

February 2001 COOK Report

Among Scaling Issues IPv6 Solves Only IP Number Problem

NATs Depend on Both IP Numbers and Routing Issues Since IPv4 and v6 Interconnect But do not Interoperate -Introducing v6 Means Running Two Networks -- 3G Makes v6 Cellular Viable

Editor's Note: The strains of scaling the Internet in the face of continued reluctance to push a switch to IPv6 produced yet another v4 versus v6 debate on the IETF mail list in mid December. We excerpt below only about 20 per cent of the total debate. We note that there are new levels of specificity of concerns indicated about the costs imposed on the v4 network of operating with address space shortage. At the end of our excerpt Fred Baker gives some fresh stats on the v4 address pool and on the extent of v6 network deployment so far.

On December 17 in a "NAT's Are Evil" discussion that had started on the IETF mail list 3 days earlier **Sean Doran** wrote: IPv6 is architecturally flawed in precisely the same way as IPv4 is; it simply has 4x the number of binary digits in the address fields and some minor cleanups which were important some years ago.

Perry Metzger: Sure. On the other hand, it has four times more bits in the address field, and at the moment, we desperately need those bits. It is costing us truly astronomical sums not to have those bits. It does not solve the routing problem. It doesn't solve world hunger either. And your point is?

(On the routing problem, the sad truth is that in spite of claims for years by you and others, there are no magic solutions to the routing problem. Blaming IPv6 for not incorporating the non-existent magic solution is rather like blaming IPv6 for not incorporating the non-existent magic cancer cure. I used to believe you and others who made vague claims about various architectures, and then I spent some time reading up on them, and I realized that none of them did particularly better.)

Bradley Dunn: I do think that there is a definite causal relationship between the address space shortage and the number of prefixes in the routing tables. People who allocate addresses .. use slow-start algorithms in their allocation policies due to the shortage of addresses. Therefore many organizations end up announcing several non-aggregatable prefixes .. .. If everyone could get "enough" addresses the first time, we'd be much closer to the ideal of one prefix per AS.

Noel Chiappa: You do have a valid point - that the address shortage is causing people to have multiple prefixes.

However, do realize that when you do an O(N) analysis on the various factors that are causing the growth in the routing table, this particular one is causing more of an O(C) factor (since most organizations with more than one prefix would have a reasonably limited number of them), or maybe O(growth-rate-of-individual-organization) [which in most cases is going to be pretty small - ISP's adding a lot of customers would be one exception]. In other words, the exponential shape of the routing table size cannot be due to an O(C) or O(average-growth-of-individual-organization) factor.

The bulk of the growth in the table has to be due to the growth in the network population as a whole, i.e. the sheer number of organizations attached to the network (which is growing at a much faster rate than any individual organization) - and in particular those which are becoming multi-homed.

It's hard to put numbers on it without knowing what %-age of sites which are already globally advertised has more that one prefix, and how fast that number is growing. However, looking at the routing table growth (it has doubled in about 3 years), and given the growth in the user community over that time, one has to expect that the growth in the user community is the driver. Yes, the point you mention will have put something of a multiplier on the table size - but it's a relativly static multiplier, I expect.

Keith Moore: if "v4" includes dealing with an increasingly severe shortage of address space (which sooner or later implies forced renumbering) and/or tying together multiple NATted networks, it's not at all clear that this takes less work than v6.

Metzger: It certainly takes more. The amount of NAT equipment out there is astonishing, and as I said at the plenary, people are starting to pay Real Money (as in millions a year) in large organizations to keep the NATs working properly. Several layers of NAT has become common, and NATs are stateful, which means they are necessarily more of a reliability problem than routers.

Multi Layers of NATs

Chiappa: This [several layers of NATs] Iam a hard time fathoming - not that I'm doubting that people do it, mind.

Metzger: Imagine a large number of companies talking to each other — the sort of situation you have when, say, you have a large clearing and settlement operation on Wall Street that has decided to speak TCP/IP to its clients. Now, imagine also that the clearing house doesn't have real IP addresses — after all, you're always told these days to use net 10 when you go to the registries and aren't going to be globally routing your nets — and that the other firms also use unregistered addresses — frequently, the same ones.

Well, you have to talk, so you use NAT.

Now, imagine that those clients have to access a service you are reselling — say, some sort of market data or specialized clearing information. That service is also delivered over TCP/IP, and also over unregistered addresses. Packets start having to traverse several address zones, just within the network obvious to the clearing organization. Now, assume that there are a couple of address zones within the client site for whatever reason, so they're using NAT internally.

Now, further assume that through various remote access schemes, you have to provide access to this mess over the "real" internet. Another layer of NAT gets added.

Are you starting to see the nightmare that has been created here? None of this is theoretical. This stuff is really happening.

Chiappa: I mean, I can understand it is a temporary thing, e.g. if one company buys another, and in gluing the networks together they temporarily leave the bought company behind a NAT, but interface it to the world via the main corporation's gateway/NAT.

Metzger: Unfortunately, multiple organizations like to talk to each other over their networks. Funny, that.

Now, you'd imagine if a Large Market Data Provider, say, went to ARIN to ask for addresses, they'd get them, but they in practice don't — they're told to use net 10 since their stuff isn't globally routed — and of course all their customers use net 10 too...

Daniel Senie: I'd read that RIPE is at least making micro-allocations available. The ability to get a few /27 allocations can REALLY help in cross-connecting corporations

which find themselves needing private interconnects. These micro-allocations are not routed globally, but rather are used to ensure unique numbers are available for private interconnects.

ARIN would do well to offer such at low cost. Or perhaps someone should gather up a bunch of allocated /24 nets which have been out there and aren't in use, and set up an interconnect registry to hand out /27s or /29s and rent them to those who need this type of private interconnect.

Chiappa: But using a NAT box adds a ration of complexity (which is always bad and a source of potential problems), and using layers of them increases the complexity, with attendant complexity costs. I have a hard time understanding why people would add that much complexity, without a darned good reason.

Metzger: They can't avoid it. They need to get their work done. They have no way of getting registered addresses. They're told to use NAT by organizations like ARIN, and so they do the only thing they can. Why do you think we're seeing huge sales of routers but somehow we haven't run out of v4 address blocks? It isn't because people are using those routers as heating equipment. It isn't because those people wouldn't prefer to get registered addresses, too.

Anyone notice how odd the growth charts look for the v4 allocation space? It is because we've already run out of addresses, folks — or at least we're acting as though we have.

I've seen as many as four layers of NAT. (That was only once, and not all the layers were within a single organization.) Two layers is routine, three less so. I'm making assumptions here, of course — you don't know what's really going on inside the other organizations. For all I know, the packets I think are coming in from the other guy's net 10 are originating behind another layer of NAT or two. Hard to tell.

It is impossible for *me* to try to figure out what is going on in such situations without a diagram. Imagine what it is like for ordinary NOC staff?

And consider that NAT boxes are stateful. When they go, they take out long lived connections, unlike dead routers, which you can simply route around.

And consider what happens when you suddenly discover that you need to re-jigger your whole nightmarish rube goldberg network because suddenly you have to make that net you never thought would have to talk to the internet show up on the net and you only have a couple of /24s you can get your hands on and have to push some of the stuff that's globally routed back into the private world somehow...

None of this is theoretical. I've seen all of this. It is astonishing how hard it has all become. As I've said, millions are being spent a year by large organizations dealing with failures and complexity attributable to NAT.

To the routing heads out there, v6 is a total failure because it doesn't solve the routing problem. I hear people like Sean Doran say "NAT is fine". That's because to the routing people and ISPs, the ugly stuff that happens at the endpoints of the networks is immaterial. ISPs get the global address blocks they want — why do they care about the rest?

Well, as things stand, we're having serious bad trouble happening at the edges of the net. NAT being a major source of operational trouble, failure and cost is not a future problem. It is here now.

You may ask why these end users aren't demanding v6. Well, they can't go to their provider and buy v6. They don't even know about v6. So they struggle along in NATworld. After all, everyone else is doing the same thing.

I mean, once you're behind a NAT box, you've got a *lot* of addresses to play with (how many, exactly, depends on how you're doing it). This is puzzling to me - what configurations are there out there that demand more address space, internally, than you already get with one layer of NAT box? Or is there some other reason I haven't figured out to have layers of address space?

Difficulties of ISPs Supporting Both v6 and v4

On Monday December 18, 2000 **RJ Atkinson:** From an operator perspective, supporting *2* IP protocols is much harder than supporting just one. If one looks around, very few NOCs on the planet today could reasonably be called "fully successful" just managing IPv4.

John Crowcoft: let's not forget that supporting IP at all is counter intuitive - the business case for providing interconnectivity to _other_ peoples services is not there, unless they are already - it requires socialist (i.e. government subsidy) to make the internet what it is now - i am betting that most telco-child ISPs love NATs and simialr technology coz they promote walled gardens (like wap etc) and lock in and all that old stuff, BUT the only way out of this dead end that is IP v4 requires either OneBigIsp to take the plunge as a way of getting more check marks on their service, or as a way of getting an even bigger wall around their garden (e.g. 3G guys might consider this:-), OR it requires some sort of socially responsible behaviour (e.g. most the 6bone is probably subsidized) to glop it all together and just make it inevitable....(this is not specially a v6 plug, just a plea for connectivity).

Atkinson: So, if one wanted IPv6 to be promoted by operators, one might spend time listening to operators and devising clever ways to make multi-homing and routing work visibly *better* with IPv6, to compensate for the much increased operational burden. Oddly enough, some folk are doing just this.

Crowcoft: Indeed - we might as well work with what we have - hey, there's a lot of stuff one can do with the code still, including just re-writing a lot of cruft in routing code - btw, [looking at] the way 3G access networks work, one ought to be able to mandate for really good aggregation - I think a lot of people forget that the exponential growth curve of the internet is not made out of homogeneous pieces - its actually a series of technology changes.

The phases from government and unviersity and dial up, and DSL and cable modem, and now mobile are all subtly different (as witness they have their own ISP cultures and own address allocation and routing headaches) - while datagram is one size fits all as a way of communicating, we need a range of address allocation and routing techniques for the different and future access and core networks....I thought most this was discussed in the whole ipng debate and was why we solicited input so widely....so now lets go (finish) coding it.

On Tuesday December 19th **Theodore Y. Ts'o** wrote: The flaw in your argument is that you're assuming that the only reason to do NAT is because of the address space problem. My concern is that it may turn out that some transport/routing people may conclude that we may also need to do NAT to solve the routing problem. In which case, we're back to where we started.

I'd feel a lot better if we could get key routing/transport people to sign some contract in blood stating that the IPV6 address is guaranteed to be invariant, and that any attempt to design boxes which muck with the IPV6 address in transit is architecturally out of bounds. That may seem to you to be obviously true, but I 10 years ago we assumed the same would be true for IPV4 addresses.

Address Space Flattening

Geoff Huston: I'm not sure that this is possible - part of the characteristics of today's Internet is that its is flattening out. The con-

cept of hierarchical connectivity with 'upstreams' and 'downstreams' is one which appears to have relied on a high marginal cost of communications services. Now as I understand the current deployment plan there are TLAs and sub TLAs, and an apparent hierarchical view of the world again.

Imagine, for a second, what the topology of the Internet would be if communications services were free. Now turn up the unit cost knob an little and do the same thought experiment.

Part of the issue we faced in the Big Internet discussions, and part of the issue we face now is the semantics of the address and the level to which these semantics are overloaded. Is an address an identification of identity? A key to absolute location? A key to relative location? An encoding of the local topology? My concern, and the reason why I'm chiming on Ted's signing blood proposal is that in looking at the structure of V6 addresses, at least the structure of the immediate deployment 1/4 or so of the addresses, we appear to have adopted an approach which is not far removed from the provider address hierarchy structure of V4 today. My lurking concern is that it is not working in the V4 routing system given the large percentage of table entries which are more specific advertisements of aggregate announcments (approx 40%) and it won't work for V6 either (using the term 'work' very loosely in terms of being able to route accurately, efficiently and with a clear scaling property).

It appears that the intended address structure and deployment structure appear to be at variance, and when that happens the temptation to alter the address in flight to suit each local region's environment may well be overwhelming. So I'd prefer to keep my blood to myself, thanks as its a contract I suspect we'll break within the operations environment. (And no I don't think that this would be a clever move, and yes, we will regret it afterwards.)

State of the IP v4 Address Pool

On Thursday December 21 **Fred Baker**: let's be fair. There was an excellent reason for NAT at the time. Postel suggested that private address spaces could be used rather than assigning precious IP Address space to networks that had no intention of attaching to the network, and NATs wound up being a way to couple that with topological address space management to try it out. We knew it was a short term hack at the time, and many of us still think that.

As Yakov is prone to point out, in a perfect world wherein all applications are client/server and address space is uniformly available, there are enough addresses around so that NATs are all we need. There are a few problems:

- the world is not perfect - all applications are not client/server - address space is not uniformly available.

Hence, NATs don't solve every problem.

The reference to IPv6 is interesting. Up until a year ago, I didn't particular push IPv6 as a solution. Reason: it wasn't in anybody's operational game plan. IPv6 had a serious chicken/egg problem - numerous people wanted to be the second to deploy it, but nobody wanted to be first, and vendors generally didn't see the point in implementing it apart from somebody waving cash to buy it. As a proposal, it solved some interesting things, like more bits in the address, better autoconfiguration, more scalable mobility, more efficient VJ Header Compression, reintroduces the end to end model so we can support non-client/server applications well, and so on.

However, being "good" isn't enough unless is it "good enough to deploy" - good enough to replace the old stuff, or good. When 3G put the proposal on the table, it became viable. At the moment, globally, we have perhaps half a dozen to a dozen commercial networks running IPv6 and upwards of 50 research networks. That's an insignificant dent in the great wide Internet, but it is not "nothing" either. We have some pretty large countries that have stated an intention to move in that direction. Now that folks have the opportunity to be second - someone else has gone first - anyone who is having trouble getting addresses from a registry is thinking seriously about IPv6.

In short, things had to get worse before they could get better. We'll see where things go, but whatever my opinions on IPv6 are (and I am on record as saying it isn't all we might have liked it to be, my voice being one among many), I am not at all convinced that it is a washout.

Tony Dal Santo: What exactly is the state of the IPv4 "address pool"? I realize there is a PERCEIVED shortage, and this is usually the main motivation for NAT. But is there a real shortage? Are "reasonable" requests for addresses being denied?

Fred Baker: The way I understand it (which could easily be out of date) is that about 45% of the address pool has been delegated, and about 25.21% is currently being advertised. The unicast address pool, what we once called the class A, class B, and class C address pools, represents 7/8 of the IP Addresses: the remainder are divided among the multicast (class D) and experimental (class E) address space.

So the bottom line is that we have delegated out a touch over half the usable unicast IP Address space. The way we are using that is, in many places, interconnecting NAT translation points - the use of private address space hides the real usage, and we have no really good way to estimate it. If we start going for non-client-server protocols - voice on IP - in a big way (and I am told that some of the world's largest telephone carriers have plans in place to convert national and international telco backbones to VoIP over the coming 3-5 years), that means that these devices will need to be addressable from outside their domains, which means those people will find themselves needing a non-NAT'd address. Implications are largely speculative, but have the option of being non-pretty.

Next question, not usually discussed, is how much of the world as yet doesn't have IP Addresses allocated to it and would like to. I think it is fair to say that the world is convinced that IP connectivity is very important. I have heard ministers of telecom from dirt-poor African countries discuss how wonderful it would be to have so much free capital laying around that they could "put a telephone into each village." Those same ministers are doing whatever it takes to ensure that their countries are on the Internet.

Unfortunately, the world is not internet-attached. Western Europe is, the US and Canada are, Australia is, Taiwan has Internet in every public library (I'm told). It comprises populations in the 1 billion person ballpark. There are some pretty large swaths of people in Eastern Europe, Asia, and Africa that are not connected and should be. If 25% of the address space is what we need for the part connected now, that tells me that I need 150% of the address space to cover everybody. If wide deployment of converged networks means that 25% was nowhere close enough for the present Internet population, then 150% is a very low guess.

So that's "what is" last I heard it from those who have the hard numbers, and "what could be". "What will be" remains anybody's guess. My crystal ball is really shiny due to excessive rubbing, and just as cloudy as ever.

February 2001 COOK Report

Napster - MP3 File Sharing Application - A Hugely Popular Bandwidth Sink

Defies Control Efforts of Network Administrators

Napster, a program written at the end of last summer by 19 year old Shawn Fanning a freshman at Northeastern University has become so popular so fast that in early March it touched off discussions on the CAIDA list and NANOG. Napster was clogging the bit pipes into and out of many University Internet connections. Some were saying that Napster was on the verge of having an impact on Internet traffic potentially as great as that of the web.

Edited versions of both discussions follow. First however let's take a quick look at what Napster actually is. According to a February 3 Janelle Brown article in Salon <http://www.salon.com/tech/feature/2000/02/03/napster/index.html> "Napster the software program — a downloadable application that lets users temporarily turn their computers into servers for the purpose of swapping MP3 files — is growing faster than anyone could have imagined. To add to the excitement, Napster the company is now embroiled in a lawsuit with the notorious Recording Industry Association of America.' The next day Scott Rosenburg <http://www.salon.com/tech/col/rose/2000/02/04/napster_swap/> commented "The brilliance of Napster is that, like the Internet itself, it lacks any center: It's just you, me and er7c, acting as individuals, sending files across the Net. The RIAA could try to shut down the central Napster directory, which lets users locate other users and their files; you can bet, though, that sooner or later someone would then come up with a more legally bulletproof version of the same service. And if the RIAA goes after the entire Napster user base, the music industry will find itself in the awkward position of suing a whole lot of its best customers. Which doesn't sound like smart business."

Not surprisingly the Recording Industry Association of America <http://www.riaa.com/piracy/press/120799.htm> expresed a different point of view when it filed suit against Napster on December 7th, 1999. "According to the complaint, Napster is similar to a giant online pirate bazaar: users log onto Napster servers and make their previously personal MP3 collections available for download by other Napster users who are logged on at the same time. Napster provides its users with all the facilities and means to engage in massive copyright infringment."

The CAIDA Discussion

On March 1, **Mark Boolootian** having complained that Napster traffic was filing up the bit pipes of the University of California Santa Cruz wrote: our survey "was based on netflow data. We took an hour snapshot of data during peak campus load and simply summed up bytes for those ports known to be used for Napster.

Vern Paxson: What I meant in particular was a question for Mark, namely: is he basing his analysis on 6699/tcp, or on a more extensive set of ports?

Boolootian: I'm sorry to say, I lied. I should have checked my sources before making any claims about how the data was collected. It turns out that the off-the-cuff analysis was not based on netflow stats. The approach taken was to block all traffic to the napster.com address space. This effectively disabled the directory service which is required for finding napster servers. A graph of the consequences of the block is at: http://noc.ucsc.edu/napster.gif

The block was in place from 16:00 until 21:30. The graph suggests that well over 50% of our outbound traffic (blue line) could be attributed to Napster. This test was performed on a Saturday, a day on which traffic flow is typically somewhat reduced from weekday traffic flow. This had the effect of slightly exaggerating the reduction.

I ran through netflow data for yesterday (Feb 29th), using ports 6688 and 6699. Outbound traffic to those two ports was 27% of total outbound campus traffic. The primary reason napster is only responsible for 27% of our outflow is, as I think I mentioned before, our local resnet police force sending cease-and-desist email to the top talkers on a daily basis.

On March 2 **Paxson**: Yowza! What is the napster.com address space? napster.com resolves for me to 208.184.216.238 and 208.49.239.246 - did you use those /24's, or something else? I'd like to do similar measurements locally ...

Boolootian: napster.com is their web site. The machine that Napster clients first talk to is server.napster.com. The Napster client connects to port 8875 on server.napster.com which returns an IP address/port number pair for one of the several Napster directory servers (those servers contain user and song info). You can see this via 'telnet server.napster.com 8875' We blocked all of 208.184.216.0/24, but it probably would have been fine to just block 208.184.216.223/32. They have changed the IP address of this box at least once - it used to be 208.184.216.222. Ran another netflow view of yesterday's outbound traffic and napster (ports 6688/6699) was 40% of total campus outbound flow. We may turn up CAR for these ports in the next day or two.

The NANOG Discussion

On NANOG on March 6 **Jeff** of the University of Tenessee at Chattanooga asked: Is there an updated list of ports used by Napster and various games (Doom, Unreal, etc)? At this point we're just collecting data points and not actively filtering anything, but of course the IANA port number list is sadly out of date with regard to these tools.

Henry Linneweh, March 7: This is an irritating problem with no update on the port list, there must be an update somewhere but it is not publicly available... And I am not finding one for Europe to see if they use different ports...

Bill Woodcock: Best port-number list I know of is the one that's included in the nmap distribution.

Joshua Goodall: Europe uses the same port numbers.

Shawn McMahon: If you're intending to filter at some point, then who cares what the ports are? Filter the ports that are using too much traffic on your network. If Napster/Doom/whatever isn't using too much bandwidth, blow it off. There's a very big difference between filtering to correct a problem, and prior restraint.

Scott McGrath: When you have limited bandwidth you need to ensure that it is used for what it is purchased for (email access to network based resources etc) and also as Napster moves MP3's you need to ensure that your facilities are not used for infringement on other's intellectual property so that you and your organization are not sued by rapacious lawyers unfortunate but there it is here is a short list in ACL format of Napster and other MP3 servers.

! Napster's Servers & Networks ! deny ip 208.49.228.0 0.0.0.255 any deny ip 208.184.216.0 0.0.0.255 any deny ip 208.49.239.240 0.0.0.15 any deny ip 208.178.175.128 0.0.0.7 any deny ip 208.178.163.56 0.0.0.7 any ! ! Other MP3 ! deny ip 202.36.148.5 0.0.0.255 any deny ip 202.36.147.16 0.0.0.255 any

Simon Lyall: In case people didn't notice these networks [206] are used by ORBS (see www.orbs.org) and people associated it with it. Perhaps some may wish to block them but it's a little pathetic for Scott to try and sneak them into other people's access lists in this way.

Scott McGrath: I am a corporate end user so the common carrier exception does not apply to me and the intellectual property demon does.

re: ORBS I got those addresses from my traffic analyis on where the MP3 traffic was coming from and I apologize to the list and it's members for not checking who the block belongs to. I have no desire to add fuel to the ORBS controversy.

Internet access has become critical to many business processes and we need to come up with mechanism to allow traffic like MP3 to fill slack in our networks however in my case I need to ensure that bandwith is available for the purpose for which it is purchased for. If Napster and it's ilk were well written applications I could create custom queues for it and discard when necessary like I do for non business related web traffic rather than my current heavy handed ACL's If anyone on the list has a alternate method I would appreciate hearing about it and my users would love me for giving them back access to Napster and other music sites.

Sean Doran: Napster traffic is principally bulk TCP transfers. Bulk TCP transfers are good for your network — much better than lots of short TCP transfers (web "mice") because they are highly sensitive to network congestion.

Moreover, where chronic congestion is a problem (i.e., your access line is full and dropping packets for long periods of time), bulk transfers back off much more quickly with a very gentle RED control law than do web mice or short SMTP sessions etc.

So, having one's supplier apply RED on your access line, and you applying RED towards the supplier, is a cheap, simple, and bandwidth-effective way of making sure that Napster traffic can consume the entire line, but not in any way worsen the performance of short TCP sessions and the like, which are not as congestion-sensitive.

I deliberately do not address intellectual property issues, nor issues of "I don't want my access line to be 100% used, even if the napster traffic behaves in a UNIX nice +20 fashion, because I pay for number of bytes moved per month".

Scott McGarth: If Napster and it's ilk were well written applications I could create custom queues for it and discard when necessary like I do for non business related web traffic rather than my current heavy handed ACL's If anyone on the list has a alternate method I would appreciate hearing about it and my users would love me for giving them back access to Napster and other music sites.

Michael Ridley from Napster.com : What would be your suggestions on making Napster a more network friendly app?

Sean Doran: 1. RFC 2001 compliance. Keep it TCP.

Michael Ridley: Agreed - I don't know of anything we would want to do that would break that stuff.

Sean Doran: 2. Although RFC 1349 is supposedly dead and the TOS octet in the RFC 791 scheme is dead too[*], it is at least good politics to set a low TOS value on the bulk transfer traffic. (If not on all traffic). Thus, routers configured to do TOS-based fancy queueing will DTRT and fewer people will accuse Napster of being a resource pig.

Michael Ridley: I believe there was some discussion about doing that, actually, although I don't know where it went (as I just do system admin stuff, not development stuff). I'll have to inquire as I haven't heard anything about it recently. Of course, the real impact would be pretty limited since I don't know that most peoples' routers really look to that header for QoS. Nevertheless...

Peter Glabavy: I think while the technology is not anywhere near perfect (hence the deprecation in most environments of ToS), if you label your traffic as bulk, then at least those who are interested can have the opportunity to prioritise based on ToS. Most will not do this, but you will gain the respect (to whatever degree) of those that do, by acknowledging that you are concerned about the resources of others. Think of it more as a politcial rather than operational acknowledgement.

Sean Doran: 3. "If we aren't network friendly, please let us know what will make us more network friendly" is a great attitude to demonstrate. Hopefully this will be appreciated by actual operators (at least the ones who don't pay per-packet/per-byte charges).

Michael Ridley: Yeah, definately we aren't trying to make peoples' lives/budgets/latencies harder than they need to be. The trick is in figuring out how much headache is the nature of the beast, how much we can help alleviate, and what the realistic compromise is.

Of course, I can only make suggestions to management, but I know that everyone here is very aware of the real life issues for the NANOG type community, and also very interested in doing what can reasonably be done.

Certainly it would be easier to queue it if it [Napster] only used one port for client->client transfer, but the problem is that there's no way to know if that port would be in use. Also, for people behind firewalls you need to tunnel in sometimes (of course, that works only a small portion of the time). I don't know what the numbers look like, but it may be the case that just queueing up the default data port would catch some large percentage of the Napster file transfering.

Scott McGarth: One of the big problems from my standpoint is that I know of no well known data port that I can queue for for web traffic from other than authorized sites I use a low priority queue which during busy times drops packets but when the usage is low all web traffic flows smoothly.

Another way would be to tag the Napster frames with QoS information so that Napster usage could be controlled with QoS policies. I really DO NOT WANT to prohibit access to Napster but I do NEED to deliver the critical services first and "fill up the corners" with entertainment services which make our end-users happy.

John Kristoff: I've been thinking and talking about this a lot with some of my peers in the University environment the last couple of weeks. Like yourself, I am trying to come up with a fairness policy and implementation based on technical limitations of IP and our network equipment.

Some of the most important thoughts/conclusions I have had are as follows: Avoid mucking with TCP and end-to-end transparency as much as possible. Let IP and RFC 2581 do its job. Managing by UDP/TCP port is unreliable. All it takes is for someone to change configuration, find the next killer app or even more fun, implement some type of "port hopping" mechanism. I'm just waiting for this to happen.

Using ToS fields *seems* to be the right approach to provide *some* of what we want. Coupled with something like WRED in our Cisco routers at the our Internet border. More sophisticated QoS mechanisms than that seems to be an effort in futility. KISS. I'm thinking a simple low priority, high priority ToS setting for packets that are

either within or outside of the local definition of fair. For example, I like the idea where all packets exceeding an average rate of X bps get tagged low priority and dropped by WRED as necessary. Be prepared to find that by using ToS bits, you break something (see http://micro.uoregon.edu/macintosh/mactcp.html).

You might be able to influence what comes into your network from the outside through these mechanisms, but you can't control them. You also can't count on any service guarantees past your administrative control. A few more good reasons to avoid anything overly complex.

It is often easier and cheaper to just add bandwidth than trying to tweak utilization and apply QoS throughout your network.

I'm having a hard time applying this fairness concept down to the user port level, if anyone has any thoughts or feedback, I'd appreciate it if you contact me (and also see my post in the unisog list).

Scott McGarth: As a member of the dreaded 'Corporate America' and an end user I need to block it for now because of 1 excessive b/w usage.

2. "Mr J. Random Engineer did you knowingly allow a service to enter your premises which allowed my clients intellectual property to be infringed yes or no?"

We need a system which allows "internet radio" for legitimate content and allows us to control bandwidth usage on end-user networks.

The common carriers have an entirely different set of needs and to keep their common carrier exemption they cannot do anything about "excessive traffic" besides if your customers want more of your product (bandwidth) more power to them assuming of course that you have a method to charge for it!

On March 8 **Eric Fair**: Any operator trying to control which applications are used on their networks has either never learned, or already forgotten, the lesson of "fsp, NASA, and OZ."

Precis: they (developers, users) can modify the application to use spread-spectrum style port numbers that constantly change throughout the life of a session, which makes it effectively impossible to filter using existing technology. Sean Doran has it right: be glad this is bulk TCP (fsp was UDP based), and turn on RED.

Scott McGarth: Well said Erik. However this is what leads many organizations to restrict/remove internet access from whole regions of the company to solve the bandwith utilization issue. So we are back to needing effective ToS/QoS methods to balance the quality of life issues with the bottom line requirements of organizations

Keeping it TCP with a well known port is a good first step because then with a custom queue you can route it with a low priority, IF the link is not busy it goes out fast this is how we handle web traffic. Sure it slows down from time to time but it ALWAYS works rather than the 'ACL the Napster servers' approach which is currently in effect. and right now Napster NEVER works.

Editor's Note: The discussion has trailed off. But Napster traffic seems uncontainable. At universities where it has been blocked, student wrote programs to redirect it to other ports.

On March 15 *Wired* carried new of AOL's moving against Gnutella, an open source successor to Napster: "Justin Frankel and Tom Pepper, founders of Nullsoft and the popular Winamp MP3 player, have been developing a file- sharing software tool which could be even more potent than Napster in letting users share MP3 files over the Net. But the "Gnutella" project may never get past its current beta stage if AOL determines the project is a threat to record labels Warner Music and EMI, which would become part of AOL this year if pending mergers are approved."

" Unlike Napster, which uses a centralized server system to let users share files, the Gnutella software was designed to create self-perpetuating networks that grow independent of one company's involvement.

After installing the program, users would connect to other "servant" computers, creating a chain of participating users. This architecture would allow for one-to-one or many-to-many connections, and makes it difficult for administrators to block the software, which occurred recently with Napster on college campuses.

"Distributed nature of servant makes it pretty damned tough for college administrators to block access to the Gnutella service.... Ability to change the port you listen on makes it even harder for those college administrators to block access," the site said before it was taken down." See http://www.wired.com/news/technology/0,1282,34978,00.html

While the AOL web site was taken down on the 15th of March a quick search of the web on March 19 showed that the Gnutilla code was already available on mirror sites. Another bandwidth eating genie is out of the bottle and pobably cannot be stuffed back in.

May 2000 COOK Report

Broadband Spread Spectrum Wireless Extends Internet Reach of ISPs & Field Research Scientists

Increased Radio Speed and Decline in Price Enabling Smaller ISPs to Compete with Cable and DSL Internet Access

NSF PI Dave Hughes Explains How Wireless Data Gathering by Environmental Scientists Will Yield Huge Increase in Network Traffic

Editor's Note: Dave Hughes is the owner of Old Colorado City Communications. See <http://www.oldcolo.com and manages, and for his NSF research projects — http://wireless.oldcolo.com>. He has been an independent networking visionary for more than twenty years. He was technical advisor to Big Sky Telegraph, linking remote schools of Montana to a University and the net in the late 1980s and early 90s. He owns his own ISP company, and uses wireless in its both up and downstream connections. Since 1995 he has been doing research for the National Science Foundation in the grass roots use of no license wireless in remote areas, including Mongolia, and beginning in the fall of 1999 no license wireless and satellite technology for biological and environmental science. We interviewed Dave on March 18, 2000. We also updated the interview in early May, 2000.

Part One
Wireless ISPs

COOK Report: It is said that there have been tremendous improvements in spread spectrum radios in the last couple of years. Are these mainly driven by better digital signal processors or by a whole lot of stuff?

Hughes: It's a combination of things. One thing that kick-started the whole rush to go from two megabit radios, (the E1 standard in Europe), to 11 megabits per second was the Harris Semiconductor 'Prism' chipset which had a much faster chipping rate and permitted the development of much faster direct sequence spread spectrum radios.

COOK Report: What is a chipping rate as opposed to a CPU rate?

Hughes: It's how fast the chipset can handle switching from one frequency to another while handling the packets of data being broadcast or received. Harris explanation is at http://www.zettweb.com/CDROMs/cdrom006/prism/

COOK Report: I see, in other words, how rapidly it can dance around.

Hughes: In a manner of speaking, yes. With direct sequence modulation, under the FCC rules, with a faster chipping rate, the FCC was able to permit the basic signal processes that allowed manufacturers across the board to jump up to a rate of 11 megabits per second from the previous 2Mbps.

COOK Report: In other words, if you get a 500% increase in your chipping rate, then you can get a 500% increase in your bandwidth?

Hughes: It's not quite that simple.

COOK Report: But it's related?

Broadband Comes to Spread Spectrum Wireless

Hughes: Yeah, it's related to the rate, and of course you realize, when you say 11 megabits per second, that you're also saying LAN. Wireless LAN, 10BaseT protocol speed. So there's no accident that everybody came roaring out at 11mbps because you will find a 10BaseT connection built into the back of all these radios. And with such a connection, you end up extending the wired LAN from inside of a building, to a wireless wide area network, or WAN. That connection can span a city, or reach across a rural county, at the same speed as if the network were all wired. Acting just like a wired network. That's a huge, huge step, because for the first time these radios operating at LAN speed can connect separate LANs without slow, or costly, wired pipes acting as a bottleneck for their performance.

Now I wasn't too impressed with Aironet radios three years ago, when I was buying radios for the first NSF project. But now they have come out with an 11 megabit per second family of radios, as did Lucent, as did Solectek and others. They all happen to be running, though, under the FCC rules for 'direct sequence' modulation, which is a different way to modulate the signal than frequency hopping. Now Breezecom, a very successful company that started in Israel and has been brought to the U.S., also has frequency-hopping radios that, for the moment, operate at less than LAN speed. They are good radios and have a large market share.

COOK Report: Okay, now do both of these "flavors" (direct sequence and frequency-hopping) fall under FCC Part 15 no license regulations?

Hughes: Oh, yes.

COOK Report: From an operational point of view why would you use direct sequence rather than frequency hoping or vice-versa?

Hughes: They differ in their ability to handle interference. Direct Sequence Spread Spectrum radios have more efficient switching by using Phase Shift Keying (PSK). They seem to be cheaper to make. And since they can operate at 11Mpbs while Frequency hopping radios are limited to 2Mbps, that is a huge advantage. The rules of the FCC to this date, have not permitted frequency hopping to reach 11 megabits per second.

COOK Report: And still do not?

Hughes: There's pressure and formal filings now at the FCC to change their rules to accommodate the new technologies. For example take Breezecom, which has a large piece of the market. They're quite successful, and we've used their radios, too. They operate at 1 megabit per second for 7 miles. The problem is they're frequency-hoppers. Given what their competition is doing, they want to come out with frequency hopping radios that operate at 11 megabits per second. And as soon as the FCC approves the changes - and there is no reason to believe they won't, then frequency hoppers can get up to the same speed, in the same frequency bands that everybody else does. But all that underscores just how bleeding edge these technologies are - the FCC is having a hard time keeping up with the rate of change.

COOK Report: What are the prices of the 11 megabit per second radios?

Hughes: They have started coming down. The early 11 Mbps Solectek Radio was $9,000 two years ago. Then out came Aironets, Lucents - competitor's that started out thousands of dollars cheaper. Right now, I can buy a PCMIA card for my laptop that

cost $165 from Aironet that operates at 11 megabits per second, half duplex. Which really means about 5-6 Mbps true throughput. Which is 3-5 times faster than a T-1. And I can buy a Teletronics 2.4Ghz, 2Mbps radio for $500 for the base unit, and only $99 for the client radio.

COOK Report: So is it safe to assume that you can buy an actual 11 megabit per second radio for under $1,000?

Hughes: Yes.

COOK Report: Way under a thousand?

Hughes: Yes. But there is a difference between buying a Base Station radio - that can handle hundreds of client radios, and a one IP/MAC address radio - for the client end. Both prices have come down rather dramatically, with client radios now below the magic 'consumer' $500 price point. Which happens to be what I had to pay for a 1200 baud Hayes Micromodem back 15 years ago.

Now let me expand on what I have been saying. This is a technical speed-up. It is one that's been coupled by more companies jumping into the game. With more learned about the software/hardware required, and with more production. Significantly, companies are now getting gobbled up by bigger companies. This is happening as these larger companies — this includes the Ciscos and Nokia's and so on —have started to get into the wireless world. And some is unlicensed, but some is licensed, like MMDS and LMDS. Which opens the door to far cheaper 'backbone' IP links.

COOK Report: Craig McCaw has huge slices of that stuff.

Hughes: Yes, because you can make radios — again, with very fast signal processors and so on — that operate in licensed spectrum, but get high rates of reliable, error corrected, reasonably secure, data communications. Up to 100Mbps.

COOK Report: And they're up into the high number of gigahertz frequency, right?

Hughes: It depends. MCI Worldcom bought up a whole lot of MMDS licenses. And they've come out of the 2.5 gigahertz to 2.7 range. That's not way up high. That's just above 2.4, which is unlicensed. And they're heading towards services that are still aimed at the corporate or business level and that get up into quite high speeds. Fixed, wireless networks.

Cisco Integrating Wireless and Wireline

Now here comes another recent development. Very, very significant. Cisco, a device-based company, manufacturer of routers par excellence ends up doing several things. It buys Aironet for $880 million. This is all in the last 60 days. They're getting into the wireless LAN market at that level. The necessary radios are down below $1,500 and with PCMIA cards down below $200 in cost. Aironet has made some changes since then, however. They have dropped the power of their radios to 30 miliwatts which is more appropriate for wireless LANs than ISPs. While you could still stick an amplifier on the antenna, doing so raises your costs. They have momentarily ceased production presumably prior to getting revved up again in their Cisco corporate guise.

COOK Report: Do the Aironet radios speak TCP/IP perfectly well?

Hughes: Oh, yes, perfectly well. Most of these spread spectrum no license radios are transparent to the protocol moving over them. But also, at the same time, Cisco itself comes out and announces an LMDS radio that'll operate 30 miles, cost $20,000 each end, but goes DS3, 45 megabits per second 30 times faster than a T-1. So there's a radio in the licensed area and that radio is out in the marketplace right now. The point is that the DS3 LMDS radio is integrated into and part of an extension of their wireline routers. Lots of integration is done in software. They're software experts. They can use the crunch capability of their router knowledge and software knowledge to drive the hardware that's already released. I have recently been briefed by them. They have some very interesting things up their sleeve. And they are not a communications service company, they are a company which makes devices - that communicate. If Cisco does things right it could represent a big paradigm shift from the Telco Model of business.

Now, remember, there are the no-licensed spread spectrum bands around 915 megahertz, and 2.4 and 5.8 gigahertz. But the FCC also came out with the UNII band.

COOK Report: That was about two years ago?

Hughes: Closer to 3 years. Now there are radios being made also in the 5.8 gigahertz UNII no-license bands that do not require spread spectrum.

COOK Report: At a 5.8 gig?

Hughes: It's the same thing. They overlap, as a matter of fact. In other words, the two band's services overlap, but that's normal in the FCC Part 15 frequency bands.

COOK Report: Aren't they severely artificially limited in their range?

Hughes: You put your finger on it, I've been screaming about this from the day it started. And you could only use one full watt of power in a UNII radio if you built a radio that offered 20 megabits of throughput. But then the FCC, always fretting about potential interference, imposed a formula for spectral density, whereby if you choose to make a radio that's only 10 megabits per second, or 5 or 1.544 - T-1 speed in the UNII bands, you must lower the power accordingly. And so every way you do it, the range of a UNII radio is not going to exceed five miles. The laws of physics still hold true. And that's okay for a lot of uses. But its lousy for anything really rural, or even school districts which are spread off across cities. Yet the UNII band was touted as having 'solved' the connectivity problem for and between public schools. My NSF studies show that no license radios need to have ranges of between 10 and 15 miles if they are to serve even 90-95% of all public school districts. Not 5 miles.

Cisco is bringing out UNII radios. They're going to be pretty low cost at the client end - perhaps from $500 to $750 at the beginning, and they're going to have other people, like Motorola manufacture the end users' radios, while they manufacture the main base engine at home, in Cisco. So my problem, which few others have expressed, is that the UNII technical radio manufacturing restrictions the FCC imposed are too limiting for rural use, where interference is not the main problem, distance of link is. The very places the President is citing as having a huge 'digital divide' - affordable rural connectivity to the Internet.

COOK Report: Clarify for me what you meant when you said $500 to $750 on the client end. In other words, a base station is, what, point to multi-point, which is more expensive and more costly than the end user device?

Hughes: Yes the base station is much more sophisticated, of course. You could even take a PCMIA card and have it talk to another PCMIA card. Period. But the minute you put a second card, two cards talking to one, then you have to have what's called a base radio, or an access point. That's essentially the generic term for it. An access point is the point at which multiple radios talk to one radio and then presumably out the back of that access radio, your traffic goes into some network by either 10BaseT or 100Mbps Ethernet. Or, as the way LANs work, by moving traffic back out via radio to another computer on the LAN, through standard hubs.

In the Aironet line, for example, is the BR500, which 11 megabits-per-second radio and operates in the 2.4-2.483 gigahertz band. It can theoretically have over a thou-

sand connections to it. Of course you are sharing its bandwidth. But that is no different from a hundred workstations on a wired LAN sharing a T-1 pipe to the Internet. Except this radio operates much faster, at least 3-5 times the throughput of a T-1. The BR500 radio itself lists at $2,400. But the end user radios, which can be a PCMIA card, can be as low as $165. And with a 'pigtail' that permits an outside antenna. One client radio talking to the point-to-multi-point radio. So there is usually a base radio - an access point - and a number of client radios, usually called bridges. They differ in price.

Now, another one of the advances that have been made is that that BR500 not only can be an access point radio, it can be configured as a relay radio. You don't have to have a different radio, as you did with the early BreezeComs. Meaning if you can't reach your destination because of non-line of sight problems or distance, you can put one BR500 at the base, one BR500 out ten miles and another one another ten miles or over the mountain. You can have 20, 30, 40 miles with one or two 'hops' using that same radio at every point. Moreover, that same radio can be logged into by password protected telnet or a web browser. They can have IP addresses in them. Even a crude router. You can use a web tool not to just configure, but also to monitor the wireless network, remotely, even if the radios are in a box high on an outside tower.

COOK Report: In other words, the newer radios come with a user-friendly GUI interface that will allow you to set it up in your own network of multiple radios?

Hughes: That's correct. I use these myself, in my own ISP operation. I originally replaced my T-1 connection to my upstream provider, 3 miles away in downtown Colorado Springs, with three Breezecom 3mbps half duplex (making them effectively the same as T-1) radios. Three radios because there is a large office building blocking a good line of sight between our building in Old Colorado City and a 14 story office building downtown Colorado Springs, where a number of ISP backbone providers are. So I set up a relay point, which, because Breezecom's Access Point/Bridge radio designs, I had to use the AP radio at the relay point. That was fine for 2 years, rain or shine. But then, as things grew and newer radios came out, I replaced the Breezecoms with 11 mbps Aironets, and simply 'configured' one of them as a relay radio.

Now, I can log into that radio that's up the street or up a tower. I don't have to climb the tower. I can log into it, I can monitor it. And when Qwest had an outage yesterday for four hours, I was able to ping all three of my radios and then finally the Cisco router, all of which responded that there was nothing wrong on my end, caused by a wireless outage. So when I called Qwest's network operation center, I knew the problem was theirs, not mine.

Then, inside my ISP offices, we have the usual wired LAN network between our servers and router. And we have a number of business customers in the same building where we are. These customers wanted higher speed connections to the net than dial up. They can't get DSL, don't like ISDN, and would have to pay, each of them, over $1,500 a month for a US West-MCI/Sprint/Quest network connections. We merely ran 10BaseT connections to them through ceiling tiles, and sell them a fast connection from $35 a month for just the link to the outside, to $55 a month for that plus full ISP service - email, web space, shell or net accounts on our server.

Now, since I want to be connected from my home-office, we have another set of radios, from Wi-Lan of Canada, 915Mhz frequency hoppers that deliver T-1 speed to me in my house. Free. And to the History Center, and its web site. And to another home-office. Then inside my house I use an older generation 'only' one mbps wireless LAN to my laptops. For use anywhere in or around my house, like out on my porch, smelling the flowers. While running my business. And connecting at the lowest speed of one mbps.

But here is another point. The Breezecoms are still in use! For a business with several workstations, in another building close by, wanted a minimum one Mbps connection. I installed the Breezecoms to link the two buildings to my net. So now I have four different radio brands connected by hubs and 10BaseT cabling to the same Wide Area Network in Old Colorado City. The moral of the story is that radios can be reused easily in different places. Lots different from yanking out wires or fiber.

And I am able, sitting at home, on a wireless LAN, to log into my Aironet radios, check the status, see how many packets it's sending, what the error rates are, or change the radio rates, set up separate wireless channels for groups of radios to operate in, within the total allocated spectrum, permitting sub or parallel nets. Check for retries, error correcting, messaging, and many other configurations. In short, all but the very lowest cost, and earlier - 3 years ago - spread spectrum radios have gotten very 'smart.' The era of the smart, self configuring radio, has arrived.

So the large companies are buying up smaller, successful wireless companies. And that's giving the whole industry credibility. And that's also what's getting the interest of Wall Street. Because when a Cisco or a Lucent and so on starts to buy up smaller companies that are successful, it's a very clear signal then they're getting into the wireless game. Whether it's licensed services on a monthly basis or whether it's unlicensed. Radically different from Telephone companies buying them up, to just put them into the tired old central office, per minute cost, circuit switched telephone business model.

This is all, of course, a relatively recent phenomenon. And while there's something in the neighborhood of 7,000 ISP's in the U.S., the fact is that hundreds of them now are delivering wireless services.

Wireless ISPs

COOK Report: Well, I heard you estimate about 300?

Hughes: That's just an estimate. There was a hard number of 180 months and months ago, but the increase is evident. It may be a thousand by now. Nobody's counting. I am watching it on a daily basis on the wireless ISP mail list. [Editor's Note: the list address is isp-wireless@isp-wireless.com and as we went to press with this article we were told that Breezecom, alone, now claims to be selling to 500 ISPs in North America].

I have been reading the wireless list for a long time, I'm seeing, on an almost daily basis, somebody new to the list saying, well, I'm an ISP and I want to offer wireless service. Or I want to start offering next month wireless ISP service from scratch. What do I need to know?

COOK Report: And they're talking a wireless service meaning connecting to an upstream wireless or allowing users to connect to them via wireless?

Hughes: No, both. Starting with users. And some of those have started out, of course, by aiming at business. And government and

Aironet BR500 11mbps Bridge, and PC C ard Direct Sequence Spread Spectrum

schools and so on.

COOK Report: It becomes another way to crack the local loop.

Hughes: Oh, definitely. It is absolutely a local loop issue. A large number of them are in pretty awkward locations. They're in rural areas where we're talking about towers for the radios.

COOK Report: And right now, we are talking local loop broadband, which as of two years ago, we weren't

Hughes: That is exactly right. Right now you have a lot of people looking at WaveSpan. UNII band. You can buy their Stratum 20 radio that goes 20 megabits per second UNII for about five miles and they maintain it can go more than that. [**Editor's Note**: Wave Span was acquired by Proxim in December 1999.] But, you can upgrade that radio to a Stratum 100 at 100 megabits per second. There are a couple of other companies which claim products in the area as well.

Well, that's significant enough that I immediately referred the San Diego Supercomputer center folks to that. Because they had an immediate, obvious need for it. The Cooperative Association for Internet Data Analysis (CAIDA) has to move its offices a mile away from the high speed network links at the Center. CAIDA needs high bandwidth for its network research. PacBell, could only talk to them maybe about a DS3, but you damn well know what telco DS3's cost. And two would be horrendous for small operation needs really fast bandwidth.

So the 100 meg radios are here. Most of them in the UNII band. Thus limited. But the FCC is also opening its mind, because there's pressure now being put on from many, many directions. In fact, there's proceedings taking place right now in which the widening of more bandwidth even at the 2.4 gigahertz, is a possibility. And there are big companies and it's controversial. Because a lot of people think the little guy's going to get squashed. But with Cisco in the game we are talking now about even bigger companies. And there are now proposals at the FCC for ultra-wide bands.

COOK Report: Ultra-wide band meaning?

Hughes: Ultra-wide band means wider than the current no-license bands that are permitted for various no-license radios. In fact, there is an ultra-wide band — you got to be real careful of the terms here — that means low power and ultra-wide to some. And no less than the engineering staff of Paul Allen's Interval Corporation, filed with the FCC. Their filing basically recommended that the band cut across all bands. And cut across television and FAA and everything else. [Editor's note: on April 21, 2000 Paul Allen announced through his Vulcan Ventures holding company that he has closed Interval Research Corp after eight years of pursuing advanced research. Approximately 30 of the staff will be offered an opportunity to join a newly formed Allen venture, which will focus purely on advanced development for Vulcan's broadband-oriented portfolio companies.]

COOK Report: Because spread spectrum is so good that it can just, it can do its thing and not interfere with others?

Hughes: They used Paul Shepherd's MIT thesis in which he mathematically demonstrated millions of radios in Manhattan, not interfering with each other, and passing hundred's of megabits per second, with radios, architecture, and of course FCC rules, to permit it. Signal processing, it isn't just spread spectrum, spread spectrum's almost a passe name now. But the idea is still, across many frequencies, sharing with other emitters and receivers, very low power - sometimes below the 'background noise' level in the area, and digitally processed.

So there are a whole series of proposals at the FCC taking place and some are being brought in by some pretty substantial companies. For example, the MMDF by MCI Worldcom, who spent $400 million getting frequencies in the 2.5 to 2.7 giga hertz range. Their Warped One would be a 1.5 megabit, a T-1, for $300 to $600 a month. But the 310 would be, called Warp 310, $40 a month for 300 KBS. MCI Worldcom, using licensed MMDS is talking about consumer level now and about direct competition with DSL and cable.

COOK Report: You have the continual slow movement of the RBOC's on DSL and, I think, the growing perception that DSL service is just going to degrade, the more DSL's you have on your copper loop from the central office. The Bell-headed telcos not doing a terrific job of rolling DSL out. The current AT&T/TCI strategy and now the AOL-Time Warner strategy for broadband access via cable TV is bogged down in controversy. Consequently, it begins to look like the shine on both DSL and cable TV as a broadband mechanism for TCP/IP is looking rather tarnished.

Hughes: You're absolutely right. Remember that DSL, is based upon the central office locations and an investment of about $150,000 to get a DSLAM into each central office. This is according to figures that I have heard directly from COVAD. And then, regardless of the overload, you start out running into another problem. As you get out toward Edge City, Suburbia, the size of the market starts dropping. And then you get into rural and small towns. Well, when it really starts dropping, then there's a crossover point where DSL is just not economic. Most of rural America will just never see it.

COOK Report: If you're looking at wireless for the lower cost issue or the issue of more or less immediate competition to cable and DSL for broadband into local loop, how important are line-of-sight issues with some of those radios?

Hughes: It's very important. However, it's interesting that Cisco, as part of their marketing, even on the frequency they're using for LMDS, maintains that they have developed technologies that limit severely the line of sight problem. In other words, diminish the problem technically.

Line of Sight and the Fresnel Zone

And one of them is called the Fresnel Zone. The Fresnel Zone is the fact that Radio waves when they travel from radio A to radio B, travel in kind of an elliptical envelope. So depending on the frequency, depending on the range, you have to have so much clearance above intervening obstacles, or from the ground at where your radio antenna is located.

I'll give a very, very specific example that I was involved with, down in the San Luis Valley. There was a tower that was down close to San Luis, which was down in a depression, with a ridge up above it. From the top of that ridge — and from San Luis to the top of the ridge is only a half a mile, maximum. But from the ridge top it's 30 miles line-of-sight to downtown Alamosa. It's a clean laser-light shot there.

The tower that we wanted to use was down on the shoulder of the ridge just above San Luis but only peeping over the top of the ridge to Alamosa, 30 miles away. By having a surveyor go up on the ridge and checking it, we found out that there was only ten feet of clearance from the top of the tower looking over the ridge line-of-sight to Alamosa. Well, it was line-of-sight from the top of the tower and it was line-of-sight down ten feet, but no further.

Fresnel Zone calculations for that radio at that range and at that frequency say you want 50 feet of clearance. And as a consequence, when, getting up on the tower, we put up a standard radio there, we found that it would not connect, because a big chunk of its power was diffused by hitting the ground before it went all the way. And yet, standing on the ridge at ground level with the same radio, with a handheld yagi antenna, with the same

power, where the ground dropped away in front of you - obviously down about 50 feet, we got a good, 30 mile connection. Consequently the only way to get around that was to go to higher power. Because we were an NSF funded experiment, we were able to get a waiver to do higher power. According to Cisco, they are now able to handle up to 49% blockage by the Fresnel Zone and still get through. So, technically, they say they have basically overcome some of the line-of-sight problems.

COOK Report: But with some of these radios, if you're talking about using them to get a broadband connection from residential homes to an ISP, you presumably would have to have that ISP radio attached to an antenna on the roof of its building. You would also quite possibly have to have an antenna on the roof of your house? In other words, if there's trees in the way, that's going to hurt also?

Hughes: You cannot generalize about things like trees. Because it all is a function of frequency, effective power and range and the nature of the obstacles. And that cannot be generalized about. So one of the fundamental realities of wireless, no matter what it is and where it is, is you must have a professional site survey before you buy and install the radios. That is unavoidable.

COOK Report: That's still the case, even with the better interfaces, and so on?

Hughes: Yes, because you're basically measuring RF signal strength Fresnel Zone obstacles. The trees, deciduous trees are not the same thing as pine trees. Wet trees are not the same thing as dry trees. Dense trees are not the same thing as sparse trees. Trees with snow hanging on the branches are not the same thing as branches in a snowstorm - when a properly sited radio will work just fine. Both through snow and rain.

COOK Report: That web site picture with the great gobs of snow on your trees, by the way, was a good visual. That picture was worth 10,000 words.

Hughes: That, of course, is it. It even blocked my little half a mile signal that's six-tenths of a watt from a Canadian Wi-Lan Radio. A good radio that is up 99.9% of the time, rain or shine over the past year.

COOK Report: That snow on the trees was like pulling down a window shade between you and the ISP.

Hughes: Yes, that's why unless you use a professional installer, you better learn a lot about how radio waves work. It's not rocket science. Certainly no harder than configuring Routers and Servers. Just different. And if distances are short, and you have good radios, you can do lots of rule of thumb work. The Aironet radio has software that permits you to walk around with a laptop and see, visually, the margin of connectivity power you have between points, with different antennas. Its the long reaches, when people try to get that last mile, out of 25, that takes skill.

Special Skills Needed by the Wireless ISPs

Hughes: Exactly right. And so you have that. So there are companies... in fact, a great deal of the discussion on the ISP list is about what does it take in distance and so on with given radios and how do you get around it? There's a guy in Texas who is superb on the subject of installations. His name is Jaime Solorza. 915-778-5966. And ask for Jaime.

The point is, they are installers. They are resellers. They are site surveyors. They do an enormous amount of business and they really know this stuff from top to bottom. And you can ask him any of these questions, but also you can even ask him for some good examples that he would know of there.

COOK Report: For example, I'm getting the impression that if you've got an ISP in a city of a couple of hundred thousand population and it's sort of semi-rural you have to have specialists. Or even if your ISP is in a city of 500 or a million population, that ISP either has to have, among its own employees or with a partnership company, someone who can do this kind of survey if it wants to advertise for wireless customers.

Hughes: Yes, on the initial set up. But it depends upon the skills of the ISP.

COOK Report: In what sense?

Hughes: Now, who are ISP's? Well, these are people who can handle Linux or Sun machines and IP and routers and so on. You've got to have that in your organization. But the ISP technical people must also learn the radio work. Siting and installation doesn't take an RF Engineering degree, but it does take some rigorous stuff, so there's lots of talk about spectrum analyzers and so on and rules of the thumb. I'm experienced enough so that I can take a look at your set-up and I can usually tell you whether it'll work or not. And if not what it would take to make it work. Know who the best friends of local ISPs can be? Ham radio operators. They may know little spread spectrum well - even though there is an entire branch of Hams, the Tucson Amateur Packet Radio group <http://www.tapr.org> who specialize in digital radio - but they sure know antennas, and wave propagation and reception. And even local laws and ordinances about placing antennas and towers or masts.

While the U.S. rule is a maximum of one watt of power at the radio (much of Europe limits it to 100 milli-watts) and then there are rules for maximum permitted antenna gain, what's called the EIRP, the effected radiated power. And then there's a whole science in antennas. And there are rules about how much total effective radiation you can have, radio power plus antenna gain. One watt at 902megahertz is going to go a lot further and through a lot more interference than at 2.4 gigahertz with the same power. A good example was the Mongolian installation that my CO-PI Dewayne Hendricks did. Because the Mongolian Engineers hadn't answered our questions about buildings in Ulaanbaatar, and what you could see from their downtown building roof - we asked for a video tape - before Dwayne arrived in Mongolia, we weren't sure whether we were going to have outside antennas on every one of those seven sites

Well, it turns out the Russians didn't put steel bars in a damn lot of those concrete buildings. And so up to about 3 kilometers, there was no requirement for an external antenna, because the 902mhz signal went through the walls to the little rubber duck antennas on the 1 watt. FreeWave serial port radios. Had there been metal in those walls, that wouldn't have been the case. And we had to guess when we shipped thousands of dollars worth of associated gear and antennas besides the radios. We would be there for just 10 days - no time for later shipment. Because we were experienced and had a little luck, we guessed right.

But this doesn't always mean that you've got to go up to the roof to get more range than a rubber duck antenna can give you. It can also mean that you have a little flat antenna just outside the window. Or a directional antenna, even inside the building, that concentrates the power and so may make the difference between success and failure. One of my NSF projects was 'Local History by Wireless' in which I put an NT web server in the computer room of an old church building, now a museum for an historical society, and connected it up 3/4 of a mile - through trees and even a big brick building - wirelessly <http://history.oldcolo.com>.

An inside ground floor antenna would not connect. An outside yagi antenna on the roof of an historically restored church would look incongruous. So I put in the attic, above the insulation blankets, but directional. An 18dbi gain yagi. It worked, and is still working 2 years later. Radios are as much an art as a science.

So there's lots of ways to adjust things, but that's part of the site survey. It's experimentation, or it is on a big network, it's really a professional thing. And the larger companies, like Solectek if you're going to buy a

whole network for a company in a town that may have six branch offices around town will send out an RF engineer, who first of all does a path analysis with software. I have a piece of software that costs $1,100. It basically takes into account the terrain, from U.S.G.S map sheets, range and then applies frequency calculations. It will give you the calculations that get you really close. From that point on, it's Reality Time, and you do what needs to be done on site.

COOK Report: So of the 300 to 500 wireless ISP's, are most of them connecting mainly a few small businesses?

Hughes: Yes, most ISPs offering wireless, offer it to businesses first. So they can learn from it, among other things. And expand from there. Until Teletronics came out with their WLAN product line, with end user radios as low as $100, the radio cost was too stiff for residential connectivity. And as these companies learned, by putting little connectors on even the PC Card radios, so external antennas could be connected, it becomes more and more possible to connect economically to home owners, and profitable for the ISP. At rates equal to or above dedicated telco services, or DSL, or cable. Bypassing them all.

There's a whole range of companies including a large one in Utah that specializes in providing high speed connections to large businesses, in direct competition with U.S. West. In many places wireless is the *only* direct competition with the telcos. I wish Congress and the FCC understood that. They flap their lips about 'competition' and about 'the digital divide' - then virtually ignore a whole emerging industry under their noses, while trying to regulate older technologies.

Wireless Business Models

COOK Report: What are the wireless ISP business models? Three flavors perhaps? ISP's that are using wireless to get to their upstream, might be one; ISP's that are using wireless to connect businesses and business offices to each other and to the ISP in place of leased lines would be another. And then is there a third where ISP's, if you want to have a radio in your home and you're just an individual user, who are beginning to do that, that's probably the most recent and the fewest?

Hughes: Yes, but that's not a very good characterization. Because, there's a totally different problem going upstream than going downstream. If you're in a big city, there's very little advantage, if you're doing really high bandwidth to be going wireless up to that upstream ISP. The real opportunity is to be found wireless in delivery that last mile or that last five miles. Or between two small towns, or suburbs.

So you can't really generalize on it. But for small towns, it's a killer, because your cost of a small town ISP is not simply getting down to your customers, your cost is also getting up to the larger city and wireless becomes very significant there — it's called microwave. And microwave is up to 100 megabits per second now. And microwave costs have come down. And so licensed microwave, which work pretty damn well when you're pushing this 50 mile range, is sometimes the method of choice to get from your town to a larger city.

COOK Report: But if you're 20 or 10 miles or something like that and you can go line of sight, then there's some other high speed radios that are pretty good.

Hughes: Right. Cisco has an LMDS radio that can do up to 30 miles, line of sight at 45 megabits. Also Jaime Solorza would be a good one to ask this question, because remember the thing I got into with Texas versus the FCC. That the FCC, as you know, has this god awful rule, involving the e-rate. Because the rule in e-rate is that the school cannot, or the library, cannot own the equipment. That is they cannot buy either the microwave equipment or even a satellite ground station (they're never cheap) or a pair of radios using e-rate funds. And provide their own connection between their building, or to the upstream ISP. Dumbest damn decision the FCC ever made.

COOK Report: That's just absurd.

Hughes: Well, it's absurd because it's expensive. And it's absurd because it's forced the schools into an annual recurring cost contractual arrangement with telcos, even though it's subsidized cost. If and when the Congress decides to kill that program, they're all going to be standing there holding the bag with their infrastructure built around that annual cost set up. The telephone companies are laughing all the way to the bank.

Now, today, for example, in the San Luis Valley is a good case, the 30-40 or more miles in the rural areas still have a $2,000 a month local loop charge for even a T-1. And with yet the radios now exist that down in the San Luis Valley, we could go on top of that ridge and we could easily be going five megabits per second true throughput for that 30 mile distance and down to that school. Doing this would cut out, totally, the $2,000 a month local loop bill for that school district. And the district would then pay only for the five megabit per second Internet bandwidth cost from the ISP, or the ISP and the school district could agree to 'choke down' their bandwidth to say, T-1, and pay less. Like maybe $1,200 a month. Now, its $2,000 a month to the telco PLUS $1,200 to the upstream ISP. Or $3,200 a month to the school district that is 40 miles away from the big city. Buy a pair, or even three, $500 2mbps radios, which incur a one time cost of less than one month of telco connectivity (which still requires that you buy a DSU to connect to them), you suddenly are down to $1,200 a month. That's the comparative economics of wireless.

And remember also that many ISP's have awakened to this, partly as a result of being approached by their customer, where the customer owns a pair of radios. The customer provides the extension from the business to the premises of the ISP, not the other way around. It's a very smart thing to do, because then you own a pair, you could take them wherever you want. And all you're doing is getting permission from the ISP or the building owner that they're in. That, of course, is what I did with Colorado Supernet 3 years ago — permission to put the antenna in there and the radio in there. Because of this 10 Base T general method, you're handing the ISP an RJ45 to connect directly into his router. He's not paying the phone company to come back to you! So the upstream Internet carrier is paying less too! Consequently, they charged me less.

So there is that absurd FCC level decision on the e-rate. You better believe I had that on my mind when I was invited to speak before the Texas Infrastructure Fund which is a Texas version of the national E-rate. And it was set up on the same principle. It comes out of rate-payers pockets and goes into a $1 billion, ten-year fund. Managed at the state level. And it was for subsidy of schools, libraries and, in the Texas case particularly, health centers.

All right. The rule was, before I made my speeches to them, that you had to have a service. If you were in Cut and Shoot, Texas, 40 miles outside of Houston, you had to use the phone company. And if that 40 miles cost $2,000 a month, that's it. You'd get the subsidy out of that fund of $2,000 a month until the end of time. And then once you're inside Cut and Shoot, there's the separate issue of how you distribute the T-1 signal, even between buildings of a company or buildings of a health center, like clinics, or the library or the school or all of the above. More monthly telephone line costs.

So I went down there to Austin, and showed the Infrastructure Fund what it could do with wireless. I said, for god's sakes, if you Texans have any brains, you'll change your rules from the way Washington and the FCC does it. They changed the rule. They may be Texans, but they aren't rubes when it comes to dollars and cents. And it's a very nice formula. The formula is, since it's based upon the phone company subsidy, essentially, if a T-1 from the ILEC or CLEC, is going to cost

$2,000 a month, they normalize it to T-1, to Cut and Shoot, or from wherever. You multiply the monthly figure by 12 and you get $24,000. And you add to it any cost of equipment that you'd have to have under the T-1. For example a DSU, CSU, or whatever.

You cowboys in Cut and Shoot may now apply all of that to an alternate means of connectivity. And it can be no license wireless, it can be licensed MSDS wireless, it could be microwave. And $24,000 will buy one hell of a lot of radio delivered bandwidth, both to the town and within it. Even relay points and everything else. And then you're only left paying for the upstream ISP. But as you know, an upstream ISP in many cases is not necessarily a commercial ISP. It may be a government. It may be a university. That's a very common thing.

As a result in Texas, any school, library, or health center, now has an alternative to continue paying monthly-telco bills with the sole future prospect of rate hikes and unending dependence on the telcos.. Would that Washington would get that smart. But then one has the feeling that the politicians are so cozy with the Telcos, they can't see over their shoulder to the Wireless Future, which is gaining on them fast.

Flexibility from Extreme Low Power and Antennas

There's one other thing I'd better say in between this. In the 2.4 gigahertz area almost everybody makes radios that only operate with 100 milli-watts of power. One-tenth of one watt, or one-tenth the authorized power by the FCC.

COOK Report: And why do they make at only a tenth of a watt?

Hughes: Because it's cheaper. It's much costlier to make a one watt radio. And they can sell the same radio in Europe, where the rule is 100 milli-watts everywhere. So they make their base radio that way because there are now many companies making amplifiers to put on line between the back of the radio and the antenna. And YDI, Young Designs, in the D.C. area, is one of the best of them. And almost every company now makes amplifiers. You put the amplifier between the radio and the antenna. If a radio puts out 100 milli-watts, you add a 1/2 watt amplifier, you are still inside FCC rules, and you can get longer range. So, again, one way to solve the problem of connectivity range or interference by trees and so on, is if you can't get through at one-tenth of one watt, you then buy a plug-in amplifier to bring you up to a full watt.

Now, I have no amplifiers on my three systems. I don't need them. I get perfect connection, because the distance is not that great. I'm not fighting anything. But if I were fighting something, I would take the $750 Young Design's 2.4 gig one watt amplifier. You can buy up to a one watt amplifier on a 2.4 gig radio. And you're still inside of the FCC rules. So amplification is becoming a very popular solution to the difficult site, the long site, the one in which the radios, even with good antennas, don't do the best job.

COOK Report: But you still can't use an amplifier to exceed the FCC limits.

Hughes: That's right. But on the other hand, these radios are so damn good that you're talking about 20 miles with one-tenth of one watt. And a pair of the lowest cost radios have been measured at almost 70 miles with line amplifiers. All within FCC rule. It works. Its legal, It's reliable. And it's secure. And it's free, free, free. Who needs Ma Bell for bandwidth?

COOK Report: I hear you.

Hughes: Well, do you know how far some have gone? When I was in San Diego, I talked to Frank Vernon, a geophysicist who works with Scrips Oceanographic Institute. Vernon is a seismologist. He monitors earthquake data. He's got a huge bunch of three-way radios that are coming back up to towers and so on. But he flat said, out loud, in front of an audience of high end scientists, without amplification, he's got one watt FreeWave radios going 100 kilometers. I keep running across some pretty long stretches. 70 miles. 50 miles. Their performance is a function, within the rules, of the height, the clear air and, if necessary, and an amplifier in there.

Software Defined Smart Radios

COOK Report: So what are they doing at the FCC?

Hughes: There's a Notice of Inquiry out and it will be discussed next week at the recently formed FCC Telecommunciations Advisory Council - TAC. - The NOI is brand new and that is on the subject of a Software Defined Radio (SDR). One where smart software controls the radio - its power, its frequency spread, and other technical characteristics.

That's what Dewayne and I recommended back in April of '98 in our *Scientific American* article. We must permit the manufacture of smart radios which set their frequencies. And set their own power levels. Keeping the power to the minimum. It's a very, very important idea. We are moving away from the idea that radios have to be dumb and fixed and made for one set of emission rules. The radios now can be smart, intelligent and self-regulating, like the Internet. But that's another way to get not only more bang for the buck, but also to minimize interference in congested areas.

If you have smart radios, built to FCC specifications so the 'self-regulation' really works, then the FCC could raise the power rules! To 10, 20 watts! Then in the really rural areas, where distance is the problem, but where interference in those bands is minimal or non-existent, they could run full power - 20 watts, with higher gain antennas. 50, 100, 250 miles. But at shorter distances and in urban areas where interference is, or can be, a problem, the radios set themselves to, say, a quarter watt. Cause that's all they need! But it has to be approved by the regulators, the FCC, which is way behind the power curve on approving these new possibilities. Maybe the recent creation of the Telecommunications Advisory Council to the FCC will help speed up change. My colleague Dewayne Hendricks is on it. And he is doing through his Dandian company on the island of Tonga, in the South Pacific. There, as they strive to link hundreds of islands, the Crown Prince sets the spectrum rules, not the FCC.

I submit the principle of 'smart radios' (and smart antennas) is a very fundamental answer to lots of 'scarce spectrum' issues nationally, and internationally, when coupled with digitally massaged data across wide bands of spectrum. George Glider talks about these things theoretically. We are doing them in the field, practically.

And remember that most of these radios also have sub channels that you can jump to. That's one way that they can it make so that everybody doesn't have to be in the same sub set. For example, in the Freewave radio, you can have 15 different settings within 902 to 928Mhz. And what does that mean? That means you can have this radio sitting here and communicating with a distant radio while you can place a radio right next to the first, operate it on the same general frequencies, and not have it interfere with the first radio.

Now, that can be set manually, of course, or they can even be set by being logged into. Now you're starting to talk about having about the little buggers scan their operating environment where they might find other Breezecoms in the area with some potential interference. Having done this, they make sure they do not interfere by setting themselves to operate in a different part of the authorized spectrum. And the FCC with the NOI (Notice of Inquiry, where it is asking 'the industry' to comment) has thrown the possibility of a world with such capabilities out there. There's going to be a debate, both technical and regulatory, because technol-

ogy that operates under these premises makes it possible change the very way that the FCC regulates spectrum.

It's not just the dumb hardware of the past, grand fathered in forever, and it's not just the fact that it's no license and it's spread spectrum. But now we're getting into the area with the software defined radios, where, if the FCC is smart, they will also shorten the life of licenses. Manufacturers must upgrade their capabilities or lose their certification for their unsold radios. For we know greater capabilities are coming along in software, radios, modulation, and antenna design. We are in an era of accelerating progress of digital radio design and operation.

COOK Report: Well, who are doing some of the offerings of the smart radios?

Hughes: Ask Jaime. I just don't memorize all the makes and models. There are at least 80 companies now, of radio manufacturers. He's in touch with damn near every one of them and he would be able to answer that real quick. Proxim just bought up one of these outfits —Wavespan but Proxim was already in this game. Research is going both ways, it's going into more powerful radios, but it's also going down to miniature radios.

Now, I put out a question all over the place — what's the smallest radio in the world? By god, I got answers. I got a reference to a Dan Withers up near Seattle and the organization is called www.worldwireless.com. I am now able to buy, a Freewave for $1,250. Oh, easy, up to 115 kilobits per second, serial. One watt. Frequency hopping. Really good radio. Very, very useful. But, whoa. I just ordered from old Dan Withers a kit, i.e., two radios, which are 56 KB, one watt, almost everything else the same characteristics. Per radio the cost is $335. And the combination is $700 for the kit, including all kinds of stuff. And you might get up as much as $500, but the fact is the pair of radios now can be bought, a serial, that'll do 56 KB or 115. For lots of uses that is plenty fast - certainly for the environmental scientists I work with, whose $3,000 data loggers put out only 9,600 baud of data.

COOK Report: At what bandwidth or what range?

Hughes: Same range, it can go up to 30 miles or 40 miles. And being 902-928 megahertz, punch through walls. Now we're talking about an end user radio inside one's house, serial. At the 902 to 928 range, frequency hopping stuff. But then there is Teletronics and their low cost 2mbps radios.

COOK Report: Well, a minute ago you said you get what you pay for. Have you tested it yet?

Hughes: You get what you pay for in the company and the support. And the total corporate follow through, ease of configuration, good documentation, best possible performance. And all those little diagnostic features.

COOK Report: So when you get the cheaper thing, you may be a little bit more on your own.

Hughes: A little bit more on your own, a little bit fewer return phone calls. But a whole mail list to ask questions on.

Customer Driven Advancement

Hughes: Yes and this leads to the concept that the end users can connect among themselves and then one of a connected group can link to an ISP. It's not all downstream.

COOK Report: Okay, because what evidence is there that, if I know the capabilities of doing these various and sundry things, and I know there is a Sprint Earthlink POP is in Trenton, New Jersey, near my house for example, I could knock on the door and ask them to let me connect? With these big national systems, there's no way in hell, with the commitment of any reasonable amount of my time they would agree to connect my radio. But what you are saying is that if the owner of a small ISP has a POP that you can reach on a reasonable basis and is aware of what can be done, even if he doesn't have a radio program yet, you can call him up and say, can I come in? You see the question I'm asking.

Hughes: That's exactly what I did with Colorado Supernet. Giant MCI would not let me do it, even though they are in the same building. They didn't have, or understand digital radios. I'm the one who went to them. They blew me off. So I went with Supernet. Now MCI keeps calling with me, pleading with me to look at their upstream prices. I blow them off now.

COOK Report: But when you go to them, do you make the argument that you should cooperate with me because even at some retail price base hook-up, I am not occupying local loop infrastructure to get into and out of your pop?

Hughes: Sure. That's part of the argument. You know what the other argument is?

COOK Report: What?

Hughes: Consider the ISP as captive to the telephone company. If the upstream ISP cooperates with you, you're essentially showing him how he may go into the wireless business by using your equipment for starters. And learn what it does. It's a cheap way for an upstream ISP to get some experience and exposure.

COOK Report: But you said there's a physical device that you can bring to him that is the receiving radio that he plugs in where?

Hughes: Into a garden variety Ethernet hub! The back side of the radio has an Ethernet port. Let me talk you through this. In my house, I have this laptop. And it's got a PC card in it that's wireless. One megabit per second. It's older. It cost $650 four years ago. I can do it now at 2mbps for $100 today. The PC card radio talks to an access point. Well, what the hell is the access point? It's nothing but a little white box into which an identical radio is plugged, like the one that goes to the PC cards. (Did you know that the much touted Apple Airport wireless is nothing more than Lucent wireless LAN cards in Apple's box?) But the only thing the box does, it turns the radio signal into an Ethernet signal. And it has a 10 Base T female connector there. That's plugged down to this little $59, five port hub. Ethernet is Ethernet.

So, between the two radios, it's radio protocol. But down to that hub, it is Ethernet. Now, coming out of the hub beneath my desk is another Ethernet that plugs into the back of this Wi-lan radio, from Canada. Which goes up to the roof to a Yagi antenna. Which is pointed towards my office. When it gets to my office wirelessly, it comes down into another radio made by the same company and set to be point to multi-point. Meaning it comes to me, but it also points down to my son David's house, so he's got a connection. And to the History Center's radio. Three of us share a T-1 connection. Could be 15 of us, in the neighborhood. Heck, one early adopter could set up an omni antenna on his house, serve 5, 10, 20 neighbors with a 2mbps or faster connection, then turn around and connect to the upstream ISP wirelessly, and split the cost 20 ways. It could be cheaper than DSL or Cable, and go where they can't or won't.

802.11 Interoperability

So I've got three different brands of radios and they're all connected and they operate at different speeds — 1 meg, T-1 and 10 megs per second. They're all normalized to an Ethernet. That's why it's not coincidental that the 10 meg is Ethernet speed. I could even go further. I could take one of these serial radios now and buy a little $40 connector that goes serial to Ethernet.

COOK Report: Go further?

Hughes: Meaning that it is not going to go Ethernet speed but it's going to go as fast as the serial port will let it. Either 56 KB or 115. But there is a demand for lower speed

radios, especially for these scientists and a lot of things, there's still a lot of serial stuff around. What I'm trying to get across here is, that you have inter connectability. It's all an extension of the LAN. And, of course, the 100 megabits is not accidental. That's 100 megabit LAN. So when you come in, you come out the back of the radio into your premises, you have 100 megabit LAN, you better have the 100 megabits, across that room to your router and so on to go upstream.

COOK Report: As long as we're talking about this topology, I had heard that it is the 802.11 standard that enables the radios to interoperate. I had heard that if I have a good connection in my house, I could connect up my neighbors to my house and then to the ISP?

Hughes: Yes. You took the words right out of my mouth, because I was about to say, and this little $150, two megabit per second Teletronics radio will talk to the $2,000 BR500 Aironet radio. The 802.11 standard interconnects them! At 2mbps.

COOK Report: So if I had a $2,000 radio in my house, I could spread out a couple of dozen 2 megabit $100 radios throughout my neighborhood.

Hughes: You got exactly what the ISP's are doing.

COOK Report: Well, which ISP's?

Hughes: A whole bunch of them. For example Jason Simonds, Midcoast Wireless. 207-563-8080. See for example: http://www.midcoast.net/wirelessfaq.html ISP for Wireless ISP's he calls himself. Now, he is an ISP. And he is doing it. As I explained much earlier in this interview, I am doing it.

An ISP in Nome, Alaska, (www.nook.net) is doing it. He is operating a dial up ISP service in Nome, connected to the net via satellite. However, he took three FreeWave radios, attached one to his dial up server in Nome. He then placed a second radio with battery and solar power, as a relay way up on a high mountain ridge 45 miles northeast of Nome. Then 25 miles beyond the relay ridge lies the village of White Mountain. He took the third radio and attached it to a dial up serviver in the village. The relay radio on the ridge can see both Nome and Whitemountain. So he delivers a commercial 56kbps connection in Whitemountain, which gets to Nome for free, and goes from there to the net via satellite. He uses the radios to extend his connection at zero additionalcost to him and at a rate of 56 kbs from Nome to White Mountain where he has paying customers. And makes a profit!

Now, a lot of this stuff is still done by hackers. Like those who ran the early Internet. And the earliest computer bulletin boards. Before AOL. Remember them? There's a woman in the wireless ISP mail list who is just so ingenious, she just drives me up the wall. But the fact is, I've been watching her ask all these questions, she acts like she's a dumb blond. But by god, she's running the thing and it's working. And she's not super high tech, but once again, where is the expertise coming from?

From the mail list. E-mail. Talking to them. And there's expertise, obviously. You've got to do a lot of learning. Eight or nine years ago it was Ethernet. And routers. It was the whole evolution of the Internet as it migrated down. 5 years ago it was the Web, and HTML coding. Where the hell was the expertise in the early 1990s? Well, it was among the hackers at the bottom and they were on the ISP list. And they were talking routers. The point is, it's still in this — while it's serious stuff for a lot of ISP's and it's real business. The fact is there's a huge amount of innovation and entrepreneurship that's taking place in communications among wannabe wireless ISP's or ISP's that want to add wireless to their operations. And they are thumbing their noses at the telcos. And setting up services where no telco dares to go. Too unprofitable for such a behemoth.

COOK Report: Well, it's this same grass roots kind of stuff which served as the training ground for all the network engineers for the commercial services.

Hughes: You're exactly right. There is really two layers to this. Cisco and the rest of the corporate world is coming down into this. For fixed wireless stuff. Either servers or like Cisco, to sell the thing and do it as a turnkey kind of set up operation. And you're having these grassroots ISP's coming up from the bottom. And they're not very small. I mean, you talk to Jaime and just ask him straight ahead what's his annual billings. I got a hunch this guy's got a hell of a growing company.

Coming from the Bottom Up

And these are in-between guys, these are not the ISP's themselves. They are resellers. But they're oh so much more than resellers. They're distributors and they're resellers and they are themselves expert in this stuff. And they go out and do the site survey. And Jaime's answering questions right and left at the same time he asks some, because this stuff is exploding in many directions.

For example, there's the issue already of throttling. This means that the ISP can adjust your radio to give you only 256 KB which you pay for. Even though your radio is capable of 10 megabits per second. This allows the ISP to price to his customers' needs. Everybody doesn't need ten megabits per second. They don't want to pay for it. It's just like any ISP. You're paying for bandwidth, right?

COOK Report: So are you saying if I'm a distributor of radios that I can make some modifications to them?

Hughes: Well, not modifications to the radio. You can use software that runs in the Linux system, for example. Or in the router. Add-ons that give you the capability to both track but also to monitor, but also to set a maximum flow rate to any given customer.

COOK Report: Okay, in other words, if I'm paying for an upstream bandwidth connection, if I have a radio that can come in and go 10 megabits and I want to put that on the guy's Ethernet, he doesn't want to give me the capability to suck up 10 megabits from his system.

Hughes: That's right. Because if the radio is capable of delivering ten megs to you, and he doesn't want to let you have the ten megs, because if he does, you will put a service on it. You will resell your connection to your neighbors. I'm doing that in Old Colorado City Communications. I could, if I wanted, spread my wireless ISP business all over 500,000 population Colorado Springs. But I have to spend half my time showing biological and environmental scientists, from those around the San Diego Supercomputer Center, to one Hispanic researcher studying frogs on the top of Mount Toro in Puerto Rico.

COOK Report: So on a small scale, you've put in your own infrastructure.

Hughes: Oh, absolutely. And have had it up for 30 months. And part of it is true wireless from me. I'm using the wireless. But things like throttling down, the ability to do that in software, there's another thing. There is already, for Linux (there must be for bigger ones), it's called EPPP. Ethernet PPP. Now you know how PPP works?

COOK Report: Yes.

Hughes: The ISP has a block from a Class B, he's got a block of numbers. And he's got X number of telephone lines, say he's got 25 lines. So he has a block of maybe 30 IP addresses. But he has 8 clients or 8 customers for every one of his 25 lines. Well, why do you have DHCP and PPP? So that when you dial in, it temporarily assigns you an IP address.

But, now there's EPPP which permits you to do that over the Ethernet. So I could have a wireless based DHCP, because right now, and I've learned about that on this list, be-

cause I basically for $5 a month, renting IP addresses from my fund of 256. And I've got them to rent, but at some point, I will run out. And so basically by having EPPP with the wireless, I can preserve my supply.

COOK Report: In other words you can take a subset of them and you can multiplex them amongst a larger group of customers.

Hughes: Yes. And with bandwidth throttling of wireless, you basically can price and measure what you do all the way up and down the line. That accounts for a whole lot of the innovation that's going on.

COOK Report: So at the grassroots, everybody, you can develop a whole mesh of interconnectedness of everybody connecting to everybody else.

Hughes: You betcha. Now, here's the telephone number and the guy you want to talk to, Michael Young, YDI, Young Designs. He's in Falls Church, Virginia. Telephone 703-237-9108. And he sells radios. He's unhappy that I bought Teletronics radios and not all from him. His radios are good, but they're a higher price. And I've been there and done that. But I have bought his amplifiers. They are better than Teletronics. Now he's measured some of these other things.

But, for me what is important is the fact that I was actually able to buy these things and get them up and get them going between two systems just lickety-split. It's almost getting to be plug-and-play. And that's from a little PC card that cost me $99 and I only paid $400 for the access point. Which could talk to multiple radios, at 2 meg per second across the room.

And then go into the Ethernet hub. You see, that's the key. The key is that by going into either a 10 Base T Ethernet series or going into the 100 megabit level, you're plugged into a purely normal networking environment. There's no magical interface. And the radios can be modulated within that framework. So that's really what's been happening with wireless ISPs.

Part 2: The NSF Field Science Research

Tachyon, Globalstar and Qualcomm

COOK Report: OK. Tell me how the NSF Field Science Project ties into all of this?

Hughes: There are two studies underway. One that is very, very significant I haven't talked much about. That's the satellite delivery of IP.

COOK Report: That ties with Tachyon into your recent San Diego Supercomputer Center meeting, doesn't it?

Hughes: Well, it does tie it in, but Tachyon is just one of them. The generalization is that, IP delivered right down to the individual by satellite is coming on eventually, with lots of services. Well, that's, okay. However, an intermediate step is being able to come down to the point in a metro area, and I don't care if it's a neighborhood center or a to business or to a school system, at such a rate of bandwidth that it is practical to distribute the signal laterally by no license wireless. Now that is one hell of a model if you think about it.

In other words, everybody thinks that the upstream ISP has got to be someplace downtown. But what happens when you are able to hook up at 2 meg up, at 2 meg down into your little ISP operation, from a satellite, or your small business operation with four offices. Or with your school system. Or with your government office. And have your IP go straight to the net from a 1 meter dish aimed at a satellite, delivering standard IP packets. But then you reach your other offices or your clients laterally by wireless that also travels at a rate of at least 2 megabits per second. Out to 1, 5, 10, 20 miles.

Tachyon so far is the only one doing this and Tachyon has its critics. But if they deliver what they promise who cares? They are just beginning to attach customers. They state that the customer ground station that talks to the satellite is only $5,000 and you can do bi-directional, true TCP/IP, at 2 megabits down and 256k up, for $2000 a month, or 300kbps down and 64k up, for $795. This will include full IP services from any spot in Europe or the Western Hemisphere. They will plug your earth station into a terrestrial wireless ISP one of which is Concentric. There is the Tachyon business model.

But you have my business model, when you extend from that base station, out 20 miles in every direction, wirelessly, and split the cost of the monthly service between 20 clients. Because then this investment is economic in every small town in America. Suddenly you don't have any phone company involved at all. And that model is really significant, because that basically becomes a real solution for the most remote towns in the U.S.

Tachyon is one of the first satellite providers which does this inside a tolerable cost envelope. It is using the SatMex5 satellite system, launched in July 1999. Therefore it should be a while before the satellite wears out. See http://www.tachyon.net for more information. The Tachyon model of course also fits the most difficult, remote, field research. One ground station, on a hill, and 10 to 100 data loggers out in every direction - all linked to it, wirelessly.

COOK Report: Is the business model emerging that Concentric will offer a family of services via Tachyon for small, remote communities who can link into the Tachyon system and then from Tachyon to Concentric to the rest of the Internet?

Hughes: Yes, that's what they are trying to do. But it also fits field scientific research, which is why I am pursuing it. Globalstar is also a good bet. Now it went down in stock price when Iridium collapsed. But what's the huge difference? Iridium was analog. Globalstar, uses Qualcomm CDMA radios. And Quaalcom is coming out with their 2.4 mbs 'HDR' wireless technology. http://www.qualcomm.com/cda/tech/hdr/

COOK Report: Well, McCaw took a look at Iridium and passed on it, but McCaw I think is investing in Globalstar, isn't he?

Hughes: Don't know. But Globalstar, has the Qualcomm CDMA spread spectrum radio. It is basically only able to deliver right now 9600 baud. But, when I made a recent cold call to Qualcomm, they were so solicitous of me, that on the same day they delivered to me the cable I needed to both charge the damn thing and use the data at the same time. Consequently out in the sticks untended it could be getting power and transmit the scientific data that we want to gather. Before they had two separate cables. You had to manually shift them to go between data use and recharging.

COOK Report: This happened at the San Diego Supercomputer Center Wireless meeting?

Hughes: Yes. With the Qualcomm radio, as they issue it right now, you get a plug, there's a plug in the bottom of the radio and you go into the recharger. Charge the radio. You unplug that thing and then you put another plug in to do a data cable, RS232 to a computer. Two different plugs. But not two different sockets. There's only a single socket for the two plugs.

COOK Report: Yeah, you have to do one or the other, but you wanted to do both.

Hughes: Exactly. And for their techs, they were doing both. But for their business model, they didn't think of that. So in their lab they had the cables.

I actually made a cold call in the morning and talked to the business section. Didn't come on real strong, all I said is that I'm

doing research for the National Science Foundation. I spent $2,600 on your stuff, including the car kit and everything else. $1,500 phone. And I said, But what I need is the cable that I understood before I bought it exists, but back at Qualcomm, not at Globalstar. I need it to hook up biological scientist's data loggers way out there where there is no cellular, no place close enough to link up with 20 mile terrestrial wireless, and of course, no telco or cableco.

And I got home at night at 9 o'clock and they had delivered to my hotel the cable that basically plugs in the bottom and has it halfway down the cable is a little plastic box that has an input to it. But much more significantly, I had a call back request to call the guy at home, the international marketing guy and when I did, he said, we'll brief you, we'll show you the next generation. I'll get that briefing in late July.

Satellite to No License Wireless Distribution

Qualcom is wireless. It is digital, not analog, as Iridium foolishly was. And it is a variation of spread spectrum. It's not free. But it's a lot less expensive than any other terrestrial solutions for really remote sites.

COOK Report: What kind of a satellite system is Tachyon using?

Hughes: They can use anybody's. They are not stuck to one. They did that intentionally. Now they are on SatMexV. They can spread their service by renting space on other birds.

COOK Report: So, in other words, they're really kind of an uplink, downlink infrastructure.

Hughes: Exactly. But bi-directional IP. That's very important. Other satellite operators are selling downlinks by satellite where the return to the Internet goes by phone lines. Of course if you are out in the wilderness with no phone this model doesn't do much good.

I talk a lot about this, because I think that's an integral part of the wireless revolution. It's wireless terrestrially, horizontally, and it's wireless vertically - to satellites. It's the combination that really makes it. I might just do this for kicks in my company. I'll get that satellite sitting on my roof here and I'll offer a separate ISP service to my neighbors, wireless. That model will work. I will get 64 KBS up and 300 KBS down for a total of about $795 a month, flat rate. $596 is Tachyon only. Tachyon plus the Concentric Internet connection is $795. Customers will normally purchase the service from Concentric — including the Tachyon ground station installation and four static IP addresses. Total throughput is measured. For example the total through put for the lowest priced service is three gigabytes per month. If customer exceeds this, he will pay 20 cents a megabyte for the extra data. The high end service allows ten gigabytes per month

COOK Report: So that pays for both the Tachyon prices and the Concentric link to the Internet.

Hughes: Yes. But what is somewhat significant on that one, it's like a telephone company demarc the Tachyon rep says, (where the phone company terminates at your premises). A demarc is where you plug into our ground station. What you do with it on the other side, it doesn't matter. The Tachyon cost is not one of these things, where if you use five computers, it costs you one thing, if it's ten, it's another. You're paying for bandwidth.

And that system — only because of the FCC — can go uplink and downlink at two megabits per second. The point is that Tachyon can offer a service at two megs now. Not what I'm getting, which is the low end 64 up and 300 KBS down. That's correct. But you see that neat little combination, because that opens the door. That opens the door not just to the U.S. That opens the door to the rest of the world.

National Environmental Observatory Network

COOK Report: Take us through a summary of the things that you saw at the San Diego meeting.

Hughes: I'll mention NEON, which everybody seems to know about. National Environmental Observatory Network. And the word 'observatory' is kind of key here, because what they're doing is gathering huge amounts of data from remote monitors.

COOK Report: So you're seeing a huge movement under foot in environmental science and in other parts of the science world to use wireless monitors.

Hughes: I would describe it as a sudden awareness that wireless is a big piece of the answer - data collection - to what they want to do.

COOK Report: And it's now economically feasible to do it and, if they get out and do it, it's going to be another huge input of bandwidth into the Internet?

Hughes: Yes. Because you must also understand the observatory concept here, the observatory means it isn't just the scientists getting data, it means that you and I can look at the damn thing. Everybody. Citizen science, said Larry Smarr. Meaning you can't afford to have high paid university researchers going out and getting all the data. This is a direct quote from him — you need to train 9-year-olds to collect data.

COOK Report: Because they want so much of it, it's so widespread?

Hughes: Because you have to. When you're talking environmental and ecological, you're dealing with a huge number of data points and all over and you've got a data collection problem that up till now has been a manual operation. Or a problem with limited resources. For example the federal government down in Puerto Rico, (the Forest Service) is out there with more expensive stuff than even the colleges use. They showed me the Sutron data collector on a stream. Sutron especially sells to governments. It's not just better equipment, it's more pricey. But it's also designed to sit out there forever on water courses, it isn't just for science engineering. It's also for monitoring flow rates, Army Corps of Engineer kinds of things. But then, it was going to a satellite, but in terms of cost effectiveness, they shut it off, it just wasn't justifiable, because it used to cost a lot to get data to the satellites.

What was NEON? The steps here is that NEON is a proposed project and they're having workshops, (this was the second) for which they bring in scientists. The focus is across disciplines from biological through environmental networking and computational. And the concept is interconnecting the scientists, the data and the databases and the visualizations and the standardization of data across disciplines. All of which requires a step up in data collection and reporting infrastructure.

What's the purpose of the meeting? The first part was for those of us who knew of technical capabilities, or in other words who knew what could be done, to be sure that the scientists sitting around there representing these various disciplines learned about the technology available to them. When they decide how the money ($100 million) should be spent, we also want to be certain that they don't think in terms of how they would have done it last year. Whether it is wireless, data bases, networking, visualization, or computation and number crunching.

Bringing the Scientists up to Wireless Speed

COOK Report: You're educating them about data gathering.

Hughes: Yes, how to gather data remotely, real time, and through the Internet right to the sensors themselves. Instead of gather-

ing the data through data loggers manually, by making visits to where the data loggers sit in the wilderness. What became very clear was that I needed to rub elbows with more of these biological scientists than I'm seeing. With just two projects right now (Puerto Rico and Wisconsin) and I needed to see what they're doing in other places and other ways.

COOK Report: You saw it there big time.

Hughes: Well, in two ways I saw it. What I was totally unprepared for, because I was just going to be an observer in this thing, was how oblivious these people were to what's available right now or has even become available in the last couple of years. They are still in the 9600 baud, RS232, coming out of a little piece of equipment, manually connected by short cables world.

And they don't need a lot more, because a lot of the data gathering is nothing but a handful of numbers. But, so I showed them, across the board, all the stuff that you can do at higher speed. And I also went to Scripps.

COOK Report: Where you saw all the earthquake sensors.

Hughes: All the earthquake stuff. They were using a very elaborate set-up, all based upon the Freewave 115 KBS radio. But they could do it, because the radio can handle the data rates that they needed. And they were using Glenayre radios also for going to a couple of points. But I knew more about radio than they did. And they were still messing around with compression as a solution to getting more bandwidth through. That is an important point, because that's getting easier to do.

Then, I took him to Tachyon. And I also found that there was Mr. Wireless for the university. There was Mr. Infrastructure. There was from up in a different campus, Mr. Data Processing guy. These were the key central guys who run the systems. And there was a discussion about what they needed to do to get the data, but they kept mentioning how expensive satellite was and they just kind of ruled it out of their minds. And I said, right here under your nose is your answer. And then, they didn't even know.

COOK Report: And Tachyon's headquarters is San Diego, yes?

Hughes: Of course. They didn't even know that Tachyon has a ground station on top of the Supercomputer Center and it goes into their network at the San Diego NAP. And, of course, if it goes into their NAP, you don't have to go out by Concentric, right? If it's a research and education application.

Data radioed from a Wisconsin lake

But Frank had not seen that data. Then there was meeting of the scientists, where I made the second presentation. To hear that you can be doing 10 megabits per second with $500 radios just blew their mental doors off. Or when Frank, sitting in the back of the room, said, "at Scripps I go 100 kilometers with a pair of three-way radios." And I said, "Boosted?"

No! Right out of the box, he replied. FCC standard regulations. Anybody can do it.

Methods of Data Collection

Okay, so they saw that, but this whole data collection stuff is based upon sensors and entering devices and data loggers and data loggers from Sutron or Campbell — Campbell is one of the big ones — these things are boxes that are smart as hell. They cost from $2, 000-10,000. They sit out there hooked to devices, like weather stations, underwater sensors, light sensors, motion, wildlife sensors. And you can have many, many devices, hundreds of them. And then they collect it.

But in, almost every case, they collect it manually into a module which then can be detached from the data collector. The module, which just a memory storage device is brought back into the lab, which is at a forward research station and dumped into a computer with the software.

COOK Report: Well, that's the old way of doing it, right?

Hughes: Yes. That's the way that's very customary. Unless it's inside a lab. But this is the field stuff. And so that's the way they're doing it.

Hughes: Now, even Campbell does sell a connection to a traditional satellite transponder service. Big cost. Expensive way of doing things. They don't think about that. It's just too damn expensive. And anyway, your Internet's got to get to the forward research station, too. Which it doesn't do down in Puerto Rico. And the research station Internet link was only 56 KB when I went up to Wisconsin.

And here's another very concrete example — and I haven't got the solution, yet, but it's very typical. Right there in Madison, Wisconsin, the main university campus sits on Lake Medora. And the University has a Center of Limnology, which is the study of great water bodies. Now they have satellites pass over and taking very costly, scientific measurements, where the colors in a photograph represent temperatures and certain chemical properties of the water and so on. The problem is calibrating actual conditions on the lake with what the satellite sees.

So the point is that they've been sending people on boats out to some 60 different points on the lake. With a graduate student and on each boat, as the satellite passes over, the student grabs the temperature and a test tube full of the water. And they come back to the lab and analyze it. They then get from those points, temperature and water composition that they use to calibrate the colors on the satellite. They have no way to 'calibrate' the satellite, with real time lake data. To do it real time, instantaneously, on a mass basis. The methodology is a labor intensive use of many people in many boats.

Obviously, wireless comes in there, because, if one power boat came roaring around the lake and dropped off a little tiny buoy that had a radio which basically broadcast that data instantaneously, on command or whatever, then that suddenly changes the nature of what they are doing. Paul Hanson, who is their chief tech, heard my pitch and ran out so fast they couldn't see straight. They now have $20,000 buoys and they bought Aironet radios. In order to see what kind of range the radio had, they put the thing

in a boat to go all the way across the lake. They ran out of water before they ran out of radio. And so all of a sudden, that's big stuff to them. Then up in northern Wisconsin, their problems are not a matter of the labor on the data points, but in getting the data from instruments, situated out in the middle of a lake, on a raft.

Now, let's talk about the levels. Ned Fetcher is a researcher in Scranton, Pennsylvania, at the university. One of the things he's been doing down there is modeling the light on the bottom of the forest floor. There's two parts to what he does. He has a data logger. And then they have these little $15, sophisticated, photosensitive light diodes. His colleague puts out twelve at a time. They have about 40 of them out there at once.

Okay, but here's the problem. They only go out to about 25 meters, 75 feet. But they have to lay a wire on the floor of the forest. But people and animals and falling branches break the wire and they lose the data.

So he said, is there any way you can get a low-cost, $50 or less, radio. And so that's when I went to this smallest radio in the world outfit and I'm having a kit shipped to me, I think it will do the job. It has a chip on it and everything else. These cheap radios become themselves data collectors. Data just goes there and into a bigger radio, back to the center and then into the data logger and be processed.

In other words twelve or more of these little radios, each with a light sensor and a pattern on the floor of the forest will communicate back to a central point no more than 300 feet away. The data is sent into a radio, where it will be passed back to he field sation.

COOK Report: So the $50 radio goes about 300 feet.

Hughes: Right. And then the radio that doesn't even have to be an Ethernet radio, it can be a serial radio, like the $300 radio. That data gets collected, then, in the data logger back at the research station which may be a mile or two away.

New Technobiology Enables New Methodologies

So that is a way to do what he wants and he was really excited about that possibility. He said it would totally transform the way that he now has to work in order to get his modeling data. The experimenter down there, a woman scientist in Puerto Rico, Jill Thompson, said they have so much bad data because of broken wires that the experiment may be useless. And all of a sudden, if they

The Coqui Frog

can put that little thing out there without wires, then all you have worry about is theradio itself being stepped on.

But you've got another advantage here. It's real time. What really turned on everybody on was the ability to look at the data in real time. Not just to get it, because it goes back not just to the research station, it also goes into the Internet at the same time. And it may have to be wireless from that point back to the university. Because the research stations are usually out in the woods. And that's exactly what we're going to be doing in Puerto Rico. With wireless we will get from the field research station back to the main university and from that out to the 'Net, so that the researcher, who may be at any university in the world can not only see this stuff, but also see if something's gone wrong. And then theresearcher can talk to a graduate student who can go out there and remove the leaf that fell over the sensor or repair whatever has to be repaired.

Now I have covered two ways of researchers getting back wireless data. One is going directly from the sensor by wireless. The other is going from a data logger by wireless, which in turn is collecting stuff from sensors, that may be very close to it. The gathered data may then go either terrestrially with a relay back to some research center, or it may go directly to a satellite. And that's where the Qualcomm phone comes in. If 9600 baud's all you need, you simply plug it in and send your data back via satellite. Unless you need to send the data constantly, you can do it periodically, at a cost of $1.50 a minute. Maybe 4 minutes a day.

But let's look at the third level where you're getting into a little bit more. The Coqui frogs. Here we want sound. Not just some sound, but quality sound. A subspecies of Coqui frog that only live on the top of mount Toro. Requiring that researchers had to climb the mountain, after dark, in the nearly perpetual rain, go into a blind, record after midnight when the frogs sing, and then come down in the morning. Very labor intensive.

Why not use radio they asked me? Sure I said. Now I have to deliver. So we've got to have enough bandwidth to make sure this is not distorted. Enough that they can't, when they record it, back at the university, lose that all the flavor. Because I noticed that when we went there, and they asked me, can we go up this mountain to do this? They had even gone back to Sony to have them optimize the microphone to improve its reception in certain frequencies, to match the frog. And so the question becomes do you have to have enough bandwidth for fidelity of sound? Therefore a low end 56kbps serial radio is probably not enough.

But then the last project I was asked to do is pretty interesting because it is full motion video from down in Puerto Rico. The guy who wants me to do it heard me when I was down there last summer talking to 40 scientists. For eight years his work has been tracking freshwater shrimp under water. And he has underwater camera doing part of the data collection. He wants me to interface the underwater camera to a high enough bandwidth to have full motion video, which can be seen, back in the upper 48 states, real time. Not

The Puerto Rican rain forest habitat of the coqui frog

just have a graduate student, as he does now, go out and manually hold a waterproof camera down there, tape record the shrimp, and send the tapes by snail mail to the distant researcher.

Now, you're talking about higher speed radios, 5 to 10 meg radios or above. And yet that's cheap to do now. Distance is not great in this case and you're not going to go to satellite with that 10 megs. But you sure as hell can go back to your research center, you may have to compress, and do all sorts of other things, but nevertheless, the radios will permit this. They will go through the forward research station to the university, wirelessly, 15 miles away. From thence into the global Internet. I've already got cameras from Axis, a Swedish company (http://www.axis.com). The cameras themselves are web servers. With individual IP addresses inside the camera. With a serial port, and an Ethernet port. The Swedes only thought you could communicate from them by either slow cell phone, or in-building Ethernet. They never thought of 2mbps wireless radios connecting them up. However, I have. We are doing it. Not full motion, only 10 jpeg frames a second. But a step above still pictures, while below full motion with sound.

Watch out UUNET, the frogs and the shrimp are coming - using your bandwidth.

And so when I was holding up these radios in front of these biological scientists, they really jumped on me after it was over, and started asking me about coming to their projects. Including for example, one up in Michigan who uses parabolic dishes to capture sounds of both insects and bird life. And they would love to have those things sitting out there all the time. And not trying to keep a tape recorder going and all that stuff.

Tiny Linux, Forest Fires and Sensors

COOK Report: All right, can we finish up with the forest fires and the sensors?

Hughes: The tiny Linux thing, including Web, is the project that I have basically asked Steve Roberts, the hacker who used a recumbent bicycle in Mountain View, to do. Remember him?

COOK Report: The Internet-connected bicycle?

Hughes: Yeah, through radio. Then, but he now, he's working right now on canoes. He's got project going down the Missouri River and the Mississippi and going up the inland waterway and then back across Canada. It'll take two years.

Well, one of the things that he was playing with there and I jumped on when I heard about it, is the concept of a very small, solar-powered, mobile data sensor capture, database and communicator. With this you grab the data on the move, not with fixed points. So he's doing that right now on a subcontract, and we call it WANDER 2000. (Wireless Acquisition of Networked Data for Environmental Research). It's a prototype, which will be done by this summer. The device will be under 15 pounds total weight. It will have a miniature Linux running in RAM and a database in it and a variety of sensors can be hooked to it.

And so whether you're on the water moving or whether you're on a trail moving, with a backpack, horse, motorized vehicle, or whatever, you can take sensors and put it into the database, which can also be a website, using Apache, if you can reach it. And communicate it by either Globestar or whatever, and by other means.

But there's a very important point here. And it was reinforced at that meeting. You've got to have reliable capture of data even if all other communication is severed, such as during a hurricane. And so the small Linux comes into play, not just that it puts up a little website and a database, but it captures reliably the data even if your communications are down for one reason or another. It gets sent when you have re-established communications. The ability to cache in a data base. That is already mastered in Data Loggers. And in institutional computers. Now we must do it when linked, wirelessly. The tool to do this is a tiny Linux server, which may be in RAM.

And I was even asked is there any way I can go out on the plains of mid America where they have firestorms, prairie fires. Can I come up with a way to communicate the temperature of the fire at ground level and the gases that are being emitted, real time, while the fire is just feet away from the radio?

Challenge? Not for wireless, that's easy. But survival has to be dealt with. Okay, the combination of these things is what drove me to this wireless outfit and I'm really excited, because first of all, that prototype called WANDER 2000 will be done by this summer and we will basically exercise it. And the full plans of that, including the wireless connectivity to it, will be basically on the website, www.worldwireless.com. But at the same time, on that other island, Whidbey Island, is where this other guy is that I tracked down yesterday.

COOK Report: Whidby Island, near Seattle?

Hughes: Whidby, I guess that's where it is. But what happened is Dan Withers with World Wireless, whom I contacted, is not only a reseller for World Wireless' very tiny, low-cost stuff, super-miniature radios that include website capability. Tiny short range ones as well as one that's only going to be $300 to go a distance. He will be presenting his engineering at a national convention of sensor manufacturers. Because at that level at which he is operating we can actually open the door to transmission from the individual sensor. A weather sensor or something and not just a big complex testing device of some kind. Pushing the radio and the wireless data collection really down to a point.

And the way you do that is you don't expect to go all the way with that little radio, but you go to a next point of aggregation and the next point of aggregation. But it's not only going up to the 'Net, it is bi-directional. And so you can have a very tiny, specialized website with the data shown on it and remote access it. Note also that this fits in with the Globe project, the observatory idea.

COOK Report: The data is collected automatically and would go into these little Databases contained in the RAM of the Linux operating system on board the radio. And the sensors, then, are feeding into a radio within a few hundred yards or something? And as part and parcel of all this you automatically fed into a remote distributed series of Linux databases that do things with it there and then feed the data back upstream?

Hughes: Yes and you used a very significant statement there.

COOK Report: Distributed?

Hughes: Distributed. Because one of the things that was even discussed at the Neon meeting was how are you going to crunch all this data? One way is centralized, terraflop computers with high bandwidth between them. The other one is the distributed process. And, see, already, there's been a model for this. I can't give you the details on it, but there was a miniature Linux that went up that was in RAM with extremely low-powered electrical draw on a space

World Wireless Communications 900 SS 56kbps low cost radio, in field protection case.

shuttle that was used for the data processing and data collection for experiments that were on board.

The Tiny Linux has three capabilities. Well, first of all, it's IP to begin with. And number two, it can have a true database. And number three, it can process. Programs can process the stuff. And then number four, it can actually also be a web point, an accessible point. I mean, two-way, not just a broadcast way or not just a capture way. And the wireless connection from it gives you the bi-directional access to it. The only thing is, you've got to control that, you can't have 10,000 people all trying to look at the damn thing at once. So you have a management problem there. But the whole idea of the observatory is that you could call up a sensor sitting in the middle of a hurricane out on the tip of Manhattan and what's it saying right now?

COOK Report: A better example would be the estuary of the Mississippi River out into the Gulf of Mexico.

Hughes: But that's going to be, I think it's going to be almost equally true in the big, urban areas. Particulates and all that. Wireless, low-cost, no license, from low to pretty high bandwidth. Sophisticated sensors. And the ability to connect to and adjust or otherwise interoperate with the censor in real time.

You have the ability to use miniature Linux as a true IP handling device, and as an IP router. It can be a router from different sensors. Remember the conviction of those who say that everything in the world is going to have an IP number in it. In the summing up after they huddled up into groups and came back with reports, on the last afternoon and I said every damn Cocqui frog's going to have an IP number. We may be putting something around his neck or embedding in its ear. Or if we're really sophisticated, we'll read it out of his DNA. DNA as IP numbers. What a gasser!

But that's, of course, fundamental to the Internet. And it's fundamental to the IP flow that's there and the wireless just permits this in places that are inconceivable. And it's going to be extremely important to biologists and environmental scientists, because their problem is dealing with data and sensing and knowing what's going on in the most remote ways, way beyond where any commercial wire line is ever going to go.

Well, in closing, everything I am doing with wireless, from remote cabins in the mountains to frogs in the rainforest, is laying down the techniques for using wireless to every human being on this planet, wherever they are. And at data rates up to full motion real time video, affordably. There is a revolution coming for the Internet, thanks to terrestrial, no license wireless, digital signal processors, smart software, IP servers, satellites and the universal connectability of the entire global net.

July 2000 COOK Report

Akamai Pushes Web Content to Edge

Content Distribution Overlay Independant of Tier 1 Backbones and Peering Is Created Using Proprietary Tools Based on DNS

Avi Freedman as Chief Architect Heads Effort to Place Servers in Networks of 1,000s of ISPs Worldwide via Incentives of Bandwith Saving & Better Response

Editor's Note: While still in college Avi Freedman founded Net Access, Philadelphia's first ISP in 1992. In the ***COOK Report*** basement in 1995, he put in a 56k-connected "POP;" however, in the mean time, he discovered DS3s, backbones, and routing. Consequently the POP never got beyond the 100-pair copper cable from the pole to our roof. Since 1995 Avi has been one of the most widely known advocates for small and mid size ISPs. In 1998 he joined AboveNet as VP of Engineering and in late 1999 he joined Akamai as Vice President and Chief Network Architect. We interviewed Avi on March 17, 2000.

COOK Report: What happened that caused you to leave you VP of Engineering position at Abovenet and go to Akamai?

Freedman: Akamai is a customer of AboveNet. And it was clear that while they had already accomplished a tremendous amount by late fall of last year, they didn't have people who had understood or had worked on big networks and knew how networks break and how they work.

COOK Report: So did you sign on with them in part to get them going in a shift from what the rest of us think of as caching to this intelligent routing that one sees discussed on their website or were they already well under way?

Freedman: They were already well under way. They already had a solution which no one else has, as far as I know. BGP tries to find the closest network by AS path. But that's not really always correct, in fact, it's often not correct. Akamai's solution uses the dynamic performance topology of the Internet to choose the best route. All performance things being equal, the algorithms may also try to push you towards a box, where the data that you want already is. If it is not there, then it will go get it. It can look like caching, but there's a lot of other stuff going on. So that's a diversion, but to answer your question, I went to talk to Akamai and was completely blown away by the number of actually smart, thoughtful people that they have and the openness of the work environment.

COOK Report: Judging by its website, Akamai has a huge network.

Freedman: We have a pretty big network and, although we don't talk about specific size, tens of gigabits per second of capacity can actually be delivered into hundreds of networks.

How Akamai Captures Talent

To answer your first question, I went there, I interviewed and it was like being back in the university. I really enjoy being in environments where there are lots of people who are smarter than I am. It means that you're guaranteed to learn anything you want, if all of you are interested in the same things. And I've been in environments where there are some smart people. But it's perhaps true that there are more smart people here than average.

COOK Report: Are a lot of them MIT?

Freedman: I think there are well over 100. For example the sales engineers at Akamai have computer science degrees from MIT.

COOK Report: What about some of the MIT TCP/IP gurus, like Dave Clark. Are any of them directly involved with the company?

Freedman: No, not directly. That's the interesting thing. Most of the MIT people who got involved were the algorithms people, who really have an important role on their own to play. To the extent people who design the algorithms were with treating networks as black boxes, they had already done a better job than everyone else. Now if you add to that the actual understanding of several specific networks and what goes on in networks overall, you get even more powerful algorithms.

Now I've also never seen a company where all the management is super excellent at what they do. You could view this as a put-down of other companies, but I really haven't seen any other company where every person in management is a so outstanding in their (and sometimes other) fields. I mean, it involves everyone, in every different phase of the company — and, they're completely open to ideas. [**Editor's Note**: At this point, sensing marketing hype I'd begin to get nervous. But I have known Avi **Freedman** since the fall of 1994 and have watched him grow and mature into one of the Internet's more widely known and respected technical figures. From my personal knowledge of Avi I believe his remarks to be absolutely sincere and genuine. Let me say also that I have no financial interests in or ties to Akamai.]

It is a revolutionary concept to me, having spent the last year essentially in California seeing mainly Silicon Valley companies. Whereas, for a company as mature as Akamai today, given that they were doing gigabit per second stuff, and just about to go public, to say to me: "We really need all this networking knowledge in the company. We want to hire you. If you accept, you're hired on the spot because we want you to hire your own staff and we want you to get them on board in the next 4 days" really got my attention.

COOK Report: I can certainly understand why. Help build some background, though, for the reader who doesn't know anything web about caching, the problems of the web as the world wide wait and things like that. Take me how we went through web caching beginning in 1996 and show me how we got from there to Akamai.

Freedman: Okay, I will disclaim that I can give you absolutely authoritative background here, but people like Peter Danzig, who invented web caching, who also work at Akamai, might be better sources. And especially, don't ask me to give exact dates.

COOK Report: Understood.

Origins of Caching

Freedman: In terms of caching, I realized sometime in the '96 range that people were doing some caching. I spent a bunch of time in Australia working with people who did networking and everyone of them did caching, because the cost of getting data there was 19 cents per megabyte, which, by the way, it still is. It is still so costly because even Telstra can't get the quantity of bandwidth, it would like to have. So caching was used for two things. Abroad, caching was used to save bandwidth and therefore to save money. In the U.S., and pretty much only the U.S. performance was the issue.

COOK Report: In other words, move your content out closer to your user, so you don't

have to go all the way across submarine cables or satellite to the U.S. to get it.

Freedman: Exactly. And people even did reverse caching, which then looks a bit like what Akamai is doing. In a place like Australia, if they had web content that they wanted available in the US they would send it to a sever located in the U.S. and serve it from there, so it didn't have constantly to be shipped across the Pacific from Australia.

COOK Report: In other words, Australian websites would be placed in the U.S.?

Freedman: In the U.S., yes. Which is somewhat like reverse caching. Which is what Inktomi is selling to people who want to compete with Akamai, but it's not what Akamai does. This gets pretty subtle.

If you look at what Akamai does, it is something that looks much more like a network than like a bunch of caches. In the U.S. in the '96 '97 time frame even, the purpose of caching was to improve performance. So, in fact, you had this box performing the caching and doing it in a way which doesn't always save bandwidth. It actually goes and pre-fetches extra data on the chance that the user might want it.

COOK Report: You might infer that it was actually using more bandwidth than they claimed to save.

Freedman: That's true, some solutions use more bandwidth than they claim to save. But it had a benefit which was that, had you infinite amounts of bandwidth, it would get extra stuff.

COOK Report: From your point of view, extra stuff more quickly. That was the purpose?

Freedman: It would attempt to have the content that the user would want a minute from now in the ready, whether or not it would actualy be asked for. It was a calculated gamble, which is okay if you have a T-3. When bandwidth costs $1,000 a megabit, you're not nearly as concerned as when bandwidth costs $60,000 a megabit. And, remember, until a year ago, bandwidth was still $10,000 a megabit in Europe. It's only now been slowing down into the couple of thousand a megabit or less.

COOK Report: Megabit per second?

Freedman: Right.. In Australia, it's megabyte. But that translates into $60,000 or $80,000 per megabit per second. Let's use Australia to understand both the value of caching and why in some ways they were the world's experts at it for a while. If bandwidth costs you 18 cents per megabyte, how much did disk space cost you?

COOK Report: Presumably, the cost of disk space is going way down.

Freedman: Oh, it's like a couple of pennies per megabyte, if that. So you might as well save everything you've ever seen on the off chance, and if 50%, you know, if anyone ever asks for it again, you win.

COOK Report: But then it gets stale, though, doesn't it?

Freedman: Yes and if bandwidth costs 19 cents a megabyte, you're actually inclined to do really strange things, like tokenize html. Now, "tokenize" is a computer science term. That means that instead of <head>, you'd use the number 1. Instead of </head>, you'd use the number 2. You convert strings of html numbers so that you can store stuff without having to run through a very CPU expensive compression, anyway. In fact, there were people that were stripping banner ads back then, because they didn't want to have to pay to bring them across and ocean to give them to the users.

COOK Report: But if they stripped banner ads in order to save some money, did the advertiser still get the value or was the advertiser bummed out?

Freedman: They simply threw the ads away. And then the other thing was that people would use what I called SECP, Social Engineering Cash Protocol. It's where you tell your users what was and you do it by looking at your cache logs. You tell you users: these are the really fast sites, and the recommended sites for tomorrow. You are confident of their speediness because you know that they are already in your cache. So it's a little bit of social engineering.

But that's the way it happened in many places outside the US.

Now in the U.S., it was a move towards faster connectivity. But how does a content provider react to all this?

A content provider says things like: what this means is we can't get tracking data. Without tracking data we cannot bill our banner advertisers and get good demographics, so we don't like you. That's what they mean.

What they say is copyright. They say, you can't store this; that's a copyright issue. Or they might say: we can't support scaling because of size constraints on number of peers, when what they really mean is, We don't want you to be a better competitor of ours.

There is also the issue for the ISP that, after a certain period of time, the content might become stale. And if you recall, there was a major U.S. backbone that actually implemented hostile caching in their network a couple of years ago and wound up screwing up a bunch of financial transactions.

COOK Report: What kind of caching?

Freedman: I call it hostile. Other people call it different things, like 'transparent'. You are doing it when you intercept users' packets. For example, in Australia, what they do is, they block port 80. Full disclosure. You cannot use our Internet connectivity to get to port 80, you must use our proxy. Full disclosure.

COOK Report: And the proxy goes through the cache first?

Freedman: The proxy is the cache.

COOK Report: And if what you want is not in the cache, does it let you out port 80?

Freedman: No, it goes gets it for you and gives it you. But you still can't talk on port 80. What the backbone was doing and what caching vendors like Inktomi are encouraging people to do is intercept port 80 packet flows and run them into a box containing their cache without telling the user what's going on. There's a euphemistic term for it, called transparent proxying or caching.. I call it hostile caching.

COOK Report: I see. But if you're doing that, also, and on the website that depends on a lot of my revenue on accurately counting hits and so on, that portion of what my site has seen elsewhere in the world, then I have no idea of what's going on. I can't capture that data. True?

Freedman: Yes.

COOK Report: And I would think people would become quite upset about that.

Freedman: Well, they are, that's why they're Akamai customers. Akamai gives the benefit of caching, it's not quite the same technology, but if you're an Akamai customer, magically, your content is served inside hundreds of locations worldwide, but you get the usage data. You know who's using what images and stuff like that.

COOK Report: The BBN Exodus thing was something entirely different?

Freedman: That was just plain peering politics. The reason that a network cannot compete with Akamai is because their peering politics prevents them from buying connectivity where needed to solve the spot bandwidth problem.

COOK Report: I am hearing about changes that would encourage business-oriented net-

works to develop along horizontal lines of connectivity and a horizontal mesh overlay of the older, more standard vertical peering practices, where everybody's either a downstream or an upstream customer of somebody else. Content delivery networks is also the newest buzzword for it. It sounds like part of what Akamai is all about.

Freedman: Let's answer that by looking at Akamai's origins. Do you know how Akamai started?

COOK Report: I'm not entirely familiar. My perception is that it started to solve the caching problem or to solve the slow World Wide Web problem.

Freedman: Right. Tim Berners Lee, who invented the Web, came to Tom Layton and Danny Lewin said, you know, we're going to have spot congestion problems.

COOK Report: At MIT surely. Well, the WC3 brought Berners Lee there, so that interaction was probably a byproduct of that, right?

Freedman: Right, and as a result, they started thinking about the problem. But they came about it from an algorithmic perspective, not with a super understanding of how networks connect. If you ask why does the Internet suck, the answer is because packets are being lost between networks. Even though there are networks that can't get packets from one side of the network to another. Most of the loss, a wide percentage, is between networks.

What Akamai Really Does

COOK Report: So in other words, what you've described it for me was the network geek solution. It was created by the network people, people like the Avi **Freedman**'s, who were doing the routing and so on. What you're saying now is that with Berners Lee and the other people at MIT you've been describing, there was a rather different way of going about it that involved algorithmic considerations, and so on. Tell me more about this.

Freedman: For something really authoritative you'd have to talk to Tom or Denny, but I will give you the picture as I understand it. There are a number of different types things that you need to do in order to get to what you could call a content distribution network. You need to be able to solve the problem of, say, ensuring freshness. Akamai has solved the problem. If you use Akamai's network to distribute your content, you will not have stale content being delivered.

COOK Report: Will you not have stale content being delivered because your server will literally be on the network itself? As opposed to sitting in a ISP content hosting web farm somewhere?

Freedman: Well, no, because Akamai doesn't have a network. Akamai doesn't have leased lines to connect its boxes together. Which is good. People say, don't you want to buy a network? And it's like, no. If I bought a network, I'd have to turn it off five times a day, because I have to make sure that I could survive if you turned off. Akamai is immune to fiber cuts.

COOK Report: Help me understand, then, how this is put together in as much detail as you possibly are able to.

Freedman: The base goal is to have servers at the edge of the Internet. What does that mean? It means ideally they're not sitting inside UUNET's hosting center, Abovenet's hosting center, or whomever else's center. Ideally, they're sitting at the ISP's, at the DSL providers, the cable head-ins. Locating in this way minimizes the problems getting to the last mile.

COOK Report: Oh, that's interesting, because for a while the idea was don't have them sitting there at the last mile, get them closer to the backbones, right?

Freedman: Some people might have said that. On the other hand that may not be particularly smart if you're trying to do content distribution.

COOK Report: I heard people saying that one way to get your stuff from one end to the other of the Internet real fast is to bring it into something like an Abovenet content hosting center, where you are very close to the backbone and your ISP has a connection that's close to the backbone, bump, you're on the backbone or OC3 or somewhere else in a millisecond, relatively speaking.

Freedman: Right. But here's another way of looking at Akamai. Akamai is an inverse hosting center, is an inverse ISP. This is the way I described it to people who knew I came from Abovenet. So, at Abovenet, the charter is take the packets and deliver them packets where they want to go with as high quality as possible. Abovenet is still the only company that I'm aware of that does it this way — your customer puts packets into your network by putting a box at an Abovenet hosting center. The way that Abovenet does that is it peers with as many people as it can. It does cold potato routing and pulls the traffic as far along its own network as possible, because its network is, on the whole, less congested than other people's. Now there is still no one else that I'm aware of that has gigabits per second that will go to every little rinky dink peering point in the world and say, hey, I'll peer with you and I'll carry all my traffic from everywhere in the world back and forth to you. No one else does that, really.

However, you still have to go get peering with people, some of whom do and some who don't want to peer with you. You have to wait for telco circuits to come in and wait to get them upgraded as you grow. And in many places, most of the ISPs that actually make up the bulk of the Internet aren't actually at peering points. And a lot of the broadband guys don't want to peer well.

COOK Report: So you are saying, then, that Abovenet does have a business practice where putting your content into its network and bringing your content close to its backbone, perhaps does make sense, in comparison at least with the alternatives.

Freedman: Right, but the alternatives there, you know, are a couple of large, well-known commercial networks. Most networks will say Oh, I handed it to them. If they can't get it cross-country, that's their problem. But Abovenet doesn't do that. I was well-trained in what a good network looks like. Now, most networks do not look like good networks. Why? because rumors to the contrary, IP bandwidth is not a commodity.

COOK Report: OK. For example, if you're on those other commercial networks and you hand the packet off to UUNET and the person that wants to receive the packet is downstream on Sprint, one potential problem is how does it bounce from UUNET over to Sprint's network, right?

Freedman: Right, exactly. So that makes sense, that's a good model. But there are tens of thousands of ISP's in the world. Some of them do not go near where Abovenet's fiber goes. And how do you get to those ISP's?

COOK Report: No. How do you get to them?

Freedman: The way Akamai does it is, it puts boxes in hundreds and then thousands of places around the world. For example, there are ways that Akamai could get a good estimate of where all the biggest ISP's were. After this is done my group would combine their traffic utilization data with BGP data and come back with the complete ranked list of how big every network is with respect to Akamai's traffic. There's some very interesting data in there that shows that a number of networks that are not nearly as big as they think are.

COOK Report: What about the MCI Cable and Wireless backbone?

Freedman: The AS3561 network was shrinking before it got divested.

COOK Report: And that AS number was

MCI's?

Freedman: Yes. I don't know the reasons but I do know people were having issues getting off that network.

COOK Report: Can you give me perhaps a little more detailed answer as to how you look at how the traffic of people who have their websites on your network, you can trace this traffic and how you monitor it?

Freedman: Networks that are well-routed want to know where their traffic goes, so that they can make sure that at least the large bits of traffic get delivered well. Abovenet is still the only multi-gigabit network that lets you see the state of all their backbone links. Whether it's problems or whether it's good or whatever. So Abovenet is actually a very good partner for Akamai. But that's a different topic.

COOK Report: OK so Akamai does go out and place its boxes all over the place, help me understand exactly what that means.

Freedman: First, networks that are using Junipers and the other big BFR's cannot get data on the new high speed router cards about where the traffic's going.

COOK Report: This is an inherent problem with the design of the router?

Freedman: It's a problem with scalability. And Junipers don't do it at all. So with Cisco, only recently and only on the low speed card, which only goes to four OC3's, can you get the data.

COOK Report: So if you've got an OC48 pipe, you can't get any data.

Freedman: That's right, OC48 cannot tell you where the traffic is going on an OC48 pipe or on a OC12 pipe on a four by OC12 card. What they've said is maybe at some point we'll give you statistical sampled data. But to the best of my knowledge, they don't have that yet. Whereas, one of the best things about looking at Akamai was, oh, here's a multi-gigabit network, in which I know exactly where the traffic's going.

COOK Report: How is that possible?

Freedman: If Akamai didn't know where the traffic was going, then it couldn't bill its customers. In order to be able to bill its customers, it has to know where the traffic is going.

The Akamai Business Model - an Intelligent Virtual Qverlay

COOK Report: So in other words, it's part of my business model. If I want to get my website on the Akamai network, I agree that I'm going to pay by some kind of traffic algorithm of data to and from my website.

Freedman: It looks just like a network. You pay X dollars per megabit on a 95th percentile addressable basis. In order to do that, we need to collect that log data, the billing data from all of our computers about to whom they served how much data. So it turns out that's the same data that tells us, if you mix it up real hard and run it through a little BGP structure of the Internet data to which networks you are sending traffic. It's actually where your traffic wants to go, where demand is. Now, if I don't have connectivity in Joe.net, I may temporarily send that into UUNET or Sprint until I get a box inside Joe.net. So that's the way of looking at it.

COOK Report: Help me understand how you get boxes inside Joe.net. How do servers get where they need to be?

Freedman: I'll just describe to you how it works. So the boxes magically appear in networks, so we start with an Akamai network that has 250 locations. In other words, we have boxes in 250 locations in a couple of hundred networks. They're the networks' locations. We just rent a rack. For a small ISP, we say, if you put our boxes in your network, here's the value proposition. You don't charge us money. We won't charge you for the boxes. The data on the boxes, will always be fresh. You'll even get streams of data, like radio, audio, video, which caches don't do.

If you're broadband, we even have customers who will be able to use our boxes as computers and serve different kinds of content to you. If you have DSL customers in Manila, now we can serve your customers there effectively. If Yahoo wants to have a bigger picture or a stream come to someone who's on broadband, well, we know which of your users are on broadband and which are not.

Now to a backbone we say, put our boxes next to every one of your major routers and if one of your users wants traffic from one of our content providers, they're going to get it from one of our servers. And why not have it be from one of servers sitting in the same city on your network versus elsewhere or versus coming in from a peering pipe? So the value proposition to a big backbone is, again, there's no caching, you don't have to redirect any of your traffic to us. You're not going to have any complaints of staleness, because the content providers are getting the data about who's using what. But you'll magically save bandwidth.

Akamai is a bandwidth-saving proposition to networks. The only people that Akamai charges are our customers, our content providers. So let's say that you sitting at home, have a computer. And you want to get to Yahoo.com, you want to get an image that's off of Yahoo.com. So the Yahoo.com web page says go to something.something.akamai.tech.net or it says something or other that hits Akamai's name servers. Your request for this URL will hit your ISP's name server. In effect if your ISP is Akamaized, you are telling its name server, hey, I want this image, which is identified by this domain name.

Then your ISP's Akamai name server asks a set of eight out of Akamai's total bunch of top level name servers saying, hey, what is the correct IP address for this domain name? (You can only specify eight at any single time, but there are really more than eight.) And Akamai's name server says, over time, the best place to go is this IP address, which is in such and such a region. Or if it doesn't know, What it says is you'll have to ask someone else.

COOK Report: What would that someone else be?

Freedman: Being another Akamai name server set. So the first set, the first time that your ISP's name server asks the question, it asks it from eight top level name servers that are always the same. And, again, it's eight of the top level name servers, but there are more than eight top level name servers in all.

COOK Report: Are these top level name servers roughly analogous to thirteen root DNS servers?

Freedman: Something like that, except they're distributed on the across different networks and they're supposed to give a good breadth of connectivity quality to far away places. The top level "eight" does not get hit that often. For all the requests to all of Akamai's customers, you're going to have to do that top level query once an hour or so or twice an hour. So it's not hit that often. Because the way the BIND works, it tries all the name servers over time and, over time, it settles on the one that seems to be the best. So then the top level Akamai name server says, Oh, for this customer, you'll have to ask these secondary eight name servers what the actual answer is. Now, of the list of eight that is returned say, six or seven of those eight, are sitting in the same place. They're sitting right in a, let's say, Netaccess or Surf.net in New Jersey or Verio in Philadelphia. Wherever Akamai thinks is a first good stab at close to where you are located. However, notice I said six of the eight. Because two of the eight are sitting somewhere completely different that's completely network topologically different and doesn't tend to go down at the same time.

COOK Report: In other words, in case there's a problem in my locality, I can go somewhere else.

Freedman: For the name servers, we're still at the name servers. Now, when you actually query them, our name servers know the state of the Internet to within a few seconds. And they know whether they should serve you or whether they should refer you, (it's called overflow) to another region. Bumping you to another region might occur because performance is bad or because they're too busy or because a machine has failed or for whatever reason.

COOK Report: In a general sense, can you explain how they know the state of the Internet to within a few seconds?

Freedman: It's a big, distributed computanetwork who has pretty bad peering to the world. But note that we have boxes in that network. So once the data gets to that box, we can serve the end users well. But the question is, how do we get our content provider's data to that box?

COOK Report: On the poorly peered network.

Freedman: Right. So if I just try to do a BGP transfer, I'd ride over the Internet. And get packet loss. Which would kill my transfer and make it really tricky. S o now I look at where is the content provider centered? Let's see the content provider.

COOK Report: In other words, you may have multiple ways of achieving the goal of getting the packets where they need to be sent and one of the first things you do is look at the various ways that that goal can be achieved.

Freedman: We can do a very precise test to tell us how our regions can get to each other and all our customers, all our content provider customers.

COOK Report: What you're describing to me (I just haven't found the right words to articulate it yet) is that kind of horizontal overlay concept way of thinking about moving traffic. One of the current buzzwords for it is Content Distribution Network.

Freedman: But we are not making networks. I mean, we are sitting on top of the networks. We're not helping networks to peer better with each other, we're just routing around them. So if, let's say the content provider is on Network C. And our user and our boxes, one set of our boxes are on Network A. If I try to go from A to C....

COOK Report: And A is the poorly peered one.

Freedman: Right. So if I try to go from A to C, I'm going to go over a poor peering pipe. And I can't change what that router's going to do. I can't tell that network, oh, to get to the our content provider, don't use C. The network is transparent to me. I can't change what it's going to do. So, but I have, let's say that Network A peers well with five out of fifty of its partners. I have boxes in those five networks and can, instead of going right to C, can reflect my traffic off of B. So I can find the correct, the unfilled, the good quality peering pipes to use to go between networks

It doesn't mean that all traffic from A to C goes well. It just, our content providers' data gets where it needs to go faster.

Because we even work with things that are not static objects, we can do live streaming and things like that. Streaming is not a question of putting the content there once and then serving it many times. You need to build a data stream back.

COOK Report: It's almost a connection-oriented stream, then, isn't it?

"Akamaizing" Content

Freedman: Well, yes, and again, there's no underlying network, because, if I had an underlying network, I'd have to worry about fiber cuts and having it fail. Instead I just rely on having lots of different networks that I can choose from at any given time. We've declared a strategy called Edge Advantage. This means we're going to give people a software interface so that they can run programs on some of our computers that are distributed out in the world. The programs allow our content providers to do, what's called Akamize their traffic and their websites. Which is an automated process. Once they do that, they'll be able to invoke banner ad insertion with just putting a simple single line of html in there principal servers. They can do this because we have banner ad insertion companies. In fact we have every important customer on the Internet. Everyone that, you know, Global Center, Abovenet and Exodus all say that they have 40% of the content. Well, we have all those customers for real.

COOK Report: I'm not quite sure what you mean by that. Are you saying some of your customers also have boxes in web farms at those other three as well as being your customers?

Freedman: All of our customers have boxes in web farms somewhere. And many of our customers have boxes in web farms at Globalcenter, Exodus, Abovenet, UUNET, AT&T, BBN.

COOK Report: I'm still trying to focus on exactly how this works. Is what you're doing, have you developed, it sounds like you've developed a means for putting intelligent devices all over the Internet that talk to each other, that figure out traffic paths and that then can figure the right traffic path to the main web server or server at some big web farm somewhere?

Freedman: That's part of it. But there is an even finer granularity to what we do. Because our content providers tell us, we know what content changes when. We can push content from the content provider to our boxes all over the world and then leave it there till it changes. Or, if the content is dynamic, we can build a communication path back from wherever the user is to the end dynamic content site. Or we can do streams over the Internet, where we reflect them in a sort of zigzag pattern, blasting through congestion to make sure that we have enough redundancy to be able to serve from, say, IDP. We can do many different things.

COOK Report:. What you're saying is, you have this rather large network of name servers but also machines that in effect can function as web servers, data storage devices and so on. Your network knows where your main customer content originating hubs are. Your network figures out how to get people who want the content to those hubs. But your network also then apparently knows how to work with the hubs or the main web farms or websites of your customers and how to bring appropriate parts of that content on an intelligently scheduled regular basis out to your own local disk storage that is scattered all over hell and gone, right?

Freedman: Bingo. You can see why I thought this was interesting.

COOK Report: Indeed. And helped along by the revolution in the crashing costs of local disk storage, too, it makes such strategy, I would think, awfully attractive.

Freedman: Yes, but remember for example, we don't have to push a whole stream. If someone's watching a movie, we don't necessarily have to go get the movie so they can look at it, we can go get chunks of the movie at a time. So, having cheap disk space is cool, but we don't rely on that. We don't have that many super disk farms around.

COOK Report: But the disk farms you do have presumably are highly flexible and highly intelligent and they can have some number of gigabytes or even terabytes of data that probably... I mean, if you look at an average disk farm that brings your customer's data out closer to the periphery,

how often does the content of this data change? Very, very frequently, I would guess.

Freedman: Well, it depends. If it's an object that doesn't change, it doesn't change. A lot of the content we serve doesn't change. A lot of the content we serve are live streams, which are one to many, and dynamic. And some of the content we serve is dynamic in nature, like in SSL, so it's one-to-one and dynamic.

COOK Report: Okay, I think Washington Post is Akamaized on your website. You have a big newspaper like Washington Post, are you saying, for example, that they would have their main web servers that they themselves are responsible for getting content changes into and out of, but that you would take their basic page format and their structural skeletal formats that presumably don't change all that often and store it all over the place.

Freedman: Right.

COOK Report: What do you do with the stuff that changes? If you have stuff on local storage and I'm requesting it through your local storage, there must be some kind of a mating process that goes on between the dynamic, brand-new, fresh data you might have to grab from somewhere else and some of the static data that you might have locally?

Freedman: Let's say you're a content provider, you have a web page. The web page lists the objects and other web pages. It's in the source code for the web page. The object is defined as, let's just say url 1. So then, each region, when our network notices that there's a new url that seems to be popular, it'll be distributed out through our network. Now, as soon as the content provider changes their web page, it'll be distributed, just so long as it's a customer.

Now, let's say that they want to change the image that's at the top of their web page, that effectively looks like it's the top of their web page. They make it a different url. Then it magically is fresh, because people stop asking for the old url. The new url gets magically distributed throughout our network. There's no freshness issue, because when the name of the object changes, we know that we have to redistribute.

That same website may have streaming going on of a live interview. Let's say CNN, because CNN is a customer of ours. We can't push that onto a hard disk because, it's live. So we then set up these virtual communication paths inside our network that enable us to take the data from CNN, distribute it in such a way that it's replicated in to each node. It's not binary. It just adaptively sends extra copies of the data so that there's lossless connectivity all the way to the end user.

COOK Report: Lossless?

Freedman: Lossless. So, that doesn't relate to storage at all. That just gets split and combined and ultimately the clean packets are served, ideally, right at the edge, a few IP hops away from where the end user is sitting because there's nothing we can do with that to make it sit on a hard disk.

COOK Report: How much of this stuff, though, does the customer want, perhaps, cached, so if I miss it, I can come back and get it 30 minutes later?

Freedman: Well, okay, but there's live streaming and there's not live streaming and they're different. But, cached is no longer an appropriate description.

COOK Report: Captured?

Freedman: Captured or put to the edge, whatever. Akamaized. The content provider has to change the name of their things so that they'll fit schema expected by the Akamai name servers.

COOK Report: In other words, if I had a live stream, I could give it one sort of Akamai URLnumber and it would be treated as a live stream. If I wanted that live stream also captured and held on disks for people to get later, I would presumably have some other code that I would assign to the data as well, is that how it works?

Freedman: That's not quite the way it works, because the live streaming versus not-live streaming is a slightly different format, actually.

COOK Report: But do you have a general principle that if I'm a customer and I have data and I want to treat it in various and sundry different ways, you can give me a menu that says here's all the things we can do with your data. The numbers on the menu might run from 1 through 75.

Freedman: We call them type codes and, yes, you can do that. And, yes, that's part of the ARL or Akamai Resource Locator. It's an extension of the URL.

But what I meant to say specifically was, if you're CNN, you can change the objects to not refer to Akamai's name servers and then you're not paying Akamai to distribute that. You have control at any time over whether a portion of your content is being distributed through Akamai.

COOK Report: What you're telling me also is that I can send out, web data, streaming data, whatever, that doesn't get distributed through Akamai, if I choose to do so and that this is data for which I don't have to pay you?

Changing Topology to Ensure Reliability

Freedman: Right. The Akamai system has never failed, but it means that were there a problem, you could just instantly stop using Akamai.

COOK Report: Or in the so-called "template" that I'm just beginning to play around with thinking about, you can stop using the horizontal mesh for distribution and go to the older more traditional vertical stream of the backbone networks and their down streams. In other words, the horizontal way of distributing, the Akamai way of distributing versus the traditional, vertical, send it up from wherever I am into the backbone, let it go across backbones and back down to the edge again. Up, across and down.

Freedman: We'd perhaps describe it a little differently. Let's just say from our perspective, it's a two-dimensional space in terms of where the packets will go. We ripped them out of their plane into a higher plane, which is the Akamai plane, shuffled them around and then drip them back into the what we view as the two-dimensional Internet.

There are other people — like Sandpiper, who is now owned by Digital Island — who actually had similar concepts, although they don't do dynamic stuff as much, and, as a result, they look more like web caching. And you say, what are the differences? The differences are we have more servers in more places and we'll always, and we know exactly where to put our servers. Our networks works, it has never failed, like most of our other competitor's networks have. And the reason is because of our good algorithms.

You know, what algorithms people look at. For example, Tom Leighton, the head of the algorithms department of the math department of MIT, so he's not a dumb person. Some of the questions on which they must focus are, will my algorithm terminate? Will it terminate and give the correct computation? As I scale, how much extra time for it take per unit input? So what's the order of the algorithm? As I scale from 100 nodes to a million, am I going to see the super-exponential breaking of the network?

This is the kind of analysis that yields a network and design that scales and doesn't break, which most of our competitors have had problems with scaling and breaking. This is the head start we have. The fact that networks want us in their network, because we have a few percent of their traffic — you know, we have Yahoo, CNN, Lycos, Mot-

ley Fool, all these websites, Barnes & Noble, Amazon — means people really want us in their network, which makes our job easier. When someone, if I wanted to start up a network to compete with Akamai, how would I convince people to put my computers in the pops of their network, give up the colo space and all that, when I have no traffic and therefore I have to pay them for it?

COOK Report: What you're getting at is this is a thing that probably can be done successfully once, period.

Freedman: Right. But of course you never say never. We're certainly not cocky. I mean, someone will come up with the next most brilliant idea.

COOK Report: Or you can do it once and probably only do it once the way you did it?

Freedman: Right. I've had people ask me, even inside Akamai who are not technical, say, okay, so we respect your opinion, Avi. Is what we're doing really rocket science? Is it really true that we have all these smart things? What I tell them is, Yes, but, at this point mostly we're writing code. The thing is, at any point, we can get 50 people together and have more brain power than anyone else in the network industry to solve any problem or react to any threat or, if the next big thing comes along. I mean, that's the cool thing about working with all these smart people. Because you never want to say no one else can do this, it's just harder.

COOK Report: You feel like you have a better critical mass of talent that you can bring to bear on any given problem.

Freedman: Right. We have 150 super-smart people in R&D. The company grew from R&D, really.

COOK Report: At one level, I see what you're doing, but at another level, I still don't see one thing in it, and I think it's terribly important and maybe you can take me through it step-by-step-by-step. And that is, I mean, you have a network, but you don't have a network. You don't own any leased lines. I think what you've been telling me in part is you've built up a critical mass of customers and tools and servers and so on where you can go to someone like a Cable & Wireless and you can tell them it's to your advantage to give us rack space for our equipment... I mean, what would you tell a big, new network that you're going into and knocking on the door for the first time, what is that you tell them?

The Economics of Distribution

Freedman: We say, right now, it's quite likely that you're paying to take packets from our customers to your customers.

COOK Report: Okay. Because it's the economic argument that I've been missing here and that's what I want you to give me.

Freedman: Or, well, there are two economics. One is the economics of our relationships with our network providers, which is sort of what I've been involuntarily promoted to be in charge of. And the other side is why it makes sense for our customers. Which is very simple, they serve more traffic to their customers, more quality traffic, their customers are happier, they buy more. Or they view more, which lets them charge money for the services they are in business to deliver.

COOK Report: In other words, hey, guy, you're a network, you've got to serve content now, as well as e-mail and mail lists, people are not going to be happy with you if you don't do a good job of serving content. Therefore, we have certain content service solutions to offer you.

Freedman: More importantly, your users are requesting content from Yahoo. There's no question about this. And I can tell you how much traffic per second I'm sending into your network. Even if I'm not sending it directly, I can tell you how much is finding its way into your network. And it's probably finding its way either via transit, which is expensive, or in just two locations, whatever your East Coast and West Coast peering is and then you have to carry it to all its destinations. It's using your long haul capacity. Not only is it using potentially your transit, or your peering, which is limited in the sense that it has to be integrated and costs some money. But it's also using your long haul capacity.

COOK Report: Okay, so you can save me money, how?

Freedman: So if you put my boxes in every one of your POP's, then you have to pay less money. For your backbone is not used as much. You can still charge your customer for delivering the packet to them if that's what your business model is, but you didn't have to pay to bring it from the East Coast to the West Coast to Ohio.

COOK Report: Well, in other words, you'll have to pay once or twice to bring it there, but you won't have to pay to do it 10,000 times a day.

Freedman: Right. And, if you have a problem in your network, we can do the dynamic routing thing, so your customer (for the content they get that's Akamaized) will still get a good experience so they don't get as angry with you if you have a problem.

COOK Report: So your argument is when you come into me and want to put boxes in my POP's, that you're going to give me access to your network and your customers. You're also going to give me access to some of your bandwidth, for which you're not going to charge me. But on the other hand, you get what is a win-win situation, because I can show you how, how in return, for that X percentage of your bandwidth that you're going to give me without charging, you can save Y amounts of money in transit charges and Z additional amounts of money in packets that you only have to get to the edge of your network and onto our boxes one time. Or you only have to get them there once a day instead of a thousand times a day or whatever.

Freedman: Right. But let's look at what happens when you hit a network that doesn't want to work with you? The reason why networks don't want to work with you is because they think they can do what you do. They don't like the fact that when one of their customers Akamaizes, their traffic goes down. Let me just tell you this, I know within a pretty good amount how much traffic every autonomous system on the Internet sends and receives. And who they're connected to.

COOK Report: And you know this because of the number of boxes you already have out there on the 'Net?

Freedman: I know this because I serve gigabits per second of traffic, I know where it goes. I have partnerships that let me get netflow data and cache logs from people, so I know where the eyeballs are. I also, on the flip side, know where the content producers are. How much they are by net, what networks they're on. Let's just say I have a network. Most networks sell to ISP's. The ISP's want us in their network, we go there. Now, they've got T-3's that they're pulling more data than they're sending. Remember a DS3 is full duplex. As much data out as in. In such a situation this network will have some unused out bound bandwidth. So I can go to that ISP and say, Hey, I'll pay you a couple of hundred bucks to get my data out.

COOK Report: I saw you do this a month ago on a mail list. So this is what all your queries were all about?.

Freedman: Yes, exactly. So if a network doesn't want to let us in their network, what's going to happen is, we're going to serve from their down stream customers and they still get no money. Their customer's utilization winds up creeping up or being symmetric, but they still aren't getting any money themselves, so they might as well let us in their network.

COOK Report: And telling them this was essentially just another way of causing

people who might be resisting working with you to realize that, if you can get the traffic out and the ISP can get the traffic out, and, if both of you can do it with highest quality and lowest cost, then they too ought to work with you?

Freedman: Yes, but some networks it's true think that they are going to buy some magic boxes, perhaps from Inktomi, and do exactly what we do. Which they can't, because they don't have boxes in other networks. They're restricted to using their own network.

COOK Report: But how all of a sudden did you grow this critical mass so damn quickly and when did you grow it?

Akamai Started Its Deployment from the Big Hosting Centers

Freedman: Yep. Well, one answer is Akamai went into all the big hosting centers. Most of them turned out to be incapable of delivering much traffic, because of congestion, bad design. You know, most hosting centers say, oh, sure, you can burst to 300 megs per second, but when you try bursting at 50, they call you up and say you're hurting their other customers.

COOK Report: So you came in there with a technical solution to a problem that initially they didn't realize they had. But they were realizing it that they had it when you came in.

Freedman: Yes. They were recognizing that, Abovenet was the only place they could burst to a gigabit on demand.

COOK Report: And some of the carrier oriented hosting centers were suffering from a certain Bell-headed mentality that caused them to be slow to understand their environment.

Freedman: And not very interested in sharing that reality with you. In other words, what date am I going to get the circuit installed on? Well, they'll give you a date, even if they know they have no intention of even putting fiber into the building by then.

So the position they were in was, the naive assumption of someone coming from a pure science background is that if someone's in the business of selling bandwidth and they want you to be able to give you the bandwidth so they can deliver it. Makes sense, right? They didn't understand that the Internet came from a model where you sell someone a pipe and hope they use none of it. This mindset sees the network as an expense and everyone wants to minimize expenses. You don't want to over-provision it, because god help you, that might encourage people to use bandwidth. Really, that is the way that many people in the Internet think.

COOK Report: And Abovenet did not think that way. For this is one reason that the web hosting centers, other than Abovenet, would have problems with people who wanted to use serious data in a bursty fashion.

Freedman: Yes. But to give credit where it is due, let me just say that Globalcenter was pretty early on with the gigabit product also. But Abovenet was like, here, take gigabits, because Abovenet realized, if people don't have gigabit speed attachments and they want to do gigabit bandwidth, they can't.

COOK Report: So in other words, some other people began to grok the Abovenet model as well.

Freedman: Right. But groking and changing are two different things. The first generation data centers were simply not of the correct scale. But note also, that AboveNet, who tried to peer with everyone, still can't get into the super edges without years of effort. It's much easier for Akamai to ship a couple of boxes into JoeNet, Omaha's local ISP, than it is to get a telco cross-connect in, or to justify the cost of a ds3 into the 8000 networks who only use a t1 or so of bandwidth.

COOK Report: I infer what you said just by my knowledge of the history of the Internet that the first generation one were those built primarily by the carriers, by the large telephone phone companies. The huge size of the Tier ones did not allow them to move as fast as you would want in the internet age to expand the data centers they had. As a result, the first generation data centers built by the first generation of internet backbone providers are the ones where you're least likely to be able to do a 100 megabits a second.

Freedman: Some reasonable inferences. In this context Akamai is going to be doing peering which for us is a very strange concept. When someone says to me: come to my peering point, I say, what AS's are there and who is likely to peer with me there? The answer to these questions determines whether it's a cost benefit for me to go there. And as far as that goes, I know how much traffic goes to each AS, so I can just do the numbers very easily and say, I need to send this much traffic to, these AS's are there, and they're the ones to which I need to send traffic.

At the same time as Akamai was deploying the big data centers, it started the outreach program that it called the Akamai Accelerated Network Program. By the end of 1999, they already had a hundred of these guys signed up. And by the end of 2000, we think we'll have 2000-3000 ISPs in the Akamai Accelerated Network Program. To put that into context,
the biggest networks may have as many as 300 peering partners.

COOK Report: A hundred of these guys being big data center guys?

Freedman: No, being ISP's. So by the time I got there full time in December, they had two things going. They had buying from networks and data centers and deploying in ISP's. What they didn't have was a set of people that were comfortable laying hands-on routers and such. And capable of talking to networks and making it clear that if they didn't give the bandwidth that they promised, that there would be detection and speedy complaining about said lack of bandwidth delivery.

COOK Report: So when you go into a network and they agree ...

Freedman: To put it bluntly, if I talk to someone and work out the architecture, they know that if they don't do it right, that I have understood exactly what we agreed to and I'm going to go back, trace the cables and get on their case. But unless the ISP's steal our hard disks or plug their boxes into our switches or try to break into our computers, then we cannot get unhappy with the ISP's.

COOK Report: Okay, give me one sentence that describes how you then come into an ISP and what you expect, what you put there. And then contrast that to what you come into a data center with, because you're saying the two are quite different.

How Being Paid for Traffic Changes the Business Model.

Freedman: With an ISP installation, we ship a switch, some computers, a terminal server, consoles, a power supply. Basically, it's all set up, except they have to type the IP addresses into the computers.

COOK Report: To make the configuration mesh into their network.

Freedman: Yes. So, obviously, they have physical access, they can steal the equipment. They could plug their computers into our switch. They could plug our computers into their switch so they can try to sniff the traffic. I mean, we've deployed in countries where they've tried to reverse engineer what we've done. They can take the memory out thinking we wouldn't notice, because each of our boxes ships with a gigabyte of memory, which is worth a few thousand dollars.

So our concerns about ISP's are only that they don't mess around with our gear, basically. We're not expecting, and they're not promising us that they're going to give us 400 megs per second connectivity. Really, they only hurt themselves if they don't give us the best connectivity they can. And they understand that. They want us to be able to steer them traffic to the users, because their users want the data that our content provider customers have.

Now our problem in dealing with hosting centers tends to be that we can't get the bandwidth we want. Because we say, we want to be able to burst 100 megabits a second. We tell them that we're going to commit to you that we'll do 10 megabits a second. Therefore, you're going to keep available 90 megabits per second. But when we find that we want to use the extra 90 megabits per second, the chances are it isn't there. Why? Because they've aggregated everyone into their routers. Because they don't understand that unless you give people the capacity you said you would for their traffic, you can't bill them for delivering the bits. And people still want to aggregate the hosting centers.

COOK Report: Well, they want to make sure that their traffic capacity in and out is fully used.

Freedman: Yes. That's the problem. Because the way IP works, you don't want to get to where you're fully utilized. You want to get to where you keep your versatility and capacity. You can't engineer to be almost fully used, because then you'll hit congestion. It has to do with the way TCP works.

People are looking at five minute averages, and therefore are not seeing the instantaneous, less than one second bursts that happen quite often. Naive people think that you can keep 90% utilization and that's a good thing to do, because you are using your resources efficiently. In reality, if you're doing 90% utilization, you're squeezing your customer's traffic, which means their flows are smaller, which means they can't send you traffic, which means you can't bill them for it, which means your revenue's lower.

A lot of people are very confused about traffic engineering and what the goal of running a network is. The goal is quantity of service, not Quality of Service.

As Dave Rand (AboveNet CTO) says, Quality of Service is where you say, if you ping me more, I'll drop his packets first. Isn't it better to be in a state where you don't have to drop packets? Where you charge everyone enough money that you can just afford to have essentially infinite bandwidth available for their use? I mean, that's the point, right? That's the point of doing networking is to be able to make happy customers and bill them and all that.

But I haven't answered the question about what is being drafted on Akamai, what network people bring. We're buying OC12's into the biggest networks that don't have good hosting centers.

COOK Report: So you are saying that you are prepared to go to some lengths to get around this problem of people in hosting centers saying: "oh, my god. You want me to be able to hold off a vast percentage of my traffic for you just in case you need it." I gather from what you said, they're rather resistant to that.

Freedman: Well, the business people agreed to it, but their technical people refused to implement it.

COOK Report: So are you saying, then, that in order to deliver the kind of Quality of Service you want to and the kind of service that gets you paid for the bits you deliver, your business model now is changing such that you are going to begin or you actually have begun to buy some of your own bit pipes?

Freedman: Yes. We are now beginning to manage routers, because that's, not because we want to, but because that's the only way we can get into some of these large networks, because you can't get in through their hosting centers, because their hosting centers are not scalable. So, fine, we'll make pipe connectivity to you. No problem.

COOK Report: And you've got to attach those big pipes of connectivity somewhere else.

Freedman: We'll just rent some space, find the closest place where the OC12 can terminate into their network, get some space and stick some boxes right next to it.

COOK Report: But now for the first time you're going to start to build and physically run your own network.

Freedman: Ah, but it's not a physical network. The point of it is, to get the OC12 in, we're not aggregating that OC12. Nor are we aggregating the network A's OC12 with the network B's OC12's and making mega-POP's. We want as fine an overlay mesh as possible. We don't want more than a percent of our connectivity going through a single point of failure.

COOK Report: I understand that. But what if you've got, say, on the face of the Earth, to choose an arbitrary number, five data centers that are giving you problems of being able to get the right connectivity into them. Let's say you're thinking of an OC3 into each of these five data centers. OK, you've got five places on the face of the Earth that your OC3's are going to begin, where are they going to be terminated?

Freedman: In the hosting center. We'll stay in the hosting center. We'll just get an OC12 into their backbone, instead of riding in their hosting center aggregation architecture. And if they don't want us in their data center doing OC12, we'll get a little hut in the back or closet somewhere and, because if the OC12 falls offline every now and then, it's no big deal, since we have more capacity to that network than we had before.

COOK Report: So you'll wind up having to pay some kind of X quantity of money to each of these five hypothetical centers that you wouldn't otherwise be paying for, but in return for the benefits that you see getting out of it, it's a worthwhile expense.

Freedman: Well, that's the other stupid thing, is that people charge more in their hosting centers than they do for their backbone connectivity.

COOK Report: So the amount of money is non-trivial?

Freedman: I'm just saying it doesn't wind up necessarily costing us more money to get a burst of OC12s than to try to get into the data centers. Because the dollars per megabit is lower on OC12s than it would be were I inside big peoples data centers.

COOK Report: Give me a specific situation.

Freedman: Joe Very Large Network wants to charge me $550 per megabit inside their hosting center. That's their bottom line price because their backbone charges them $500 a megabit, because they want to make profit centers. But they'll sell to an ISP at $400 per megabit.

COOK Report: Why do they do that?

Freedman: They try to sell their hosting center to content providers as a premium space and power and one, theoretically, some redundancy and scalability of facilities. But that's a different question.

COOK Report: There are extra charges for the physical facilities, but they make the bandwidth a little bit cheaper?

Freedman: No, they make the bandwidth more expensive at hosting centers. They charge $500, $600 a megabit per second in the hosting centers. But remember that these older hosting centers are each connected by only two OC3's and they have a hundred customers in them. Now how could they possibly give Akamai 300 megabits per second in connectivity if they only have 300 megabits per second into their hosting cen-

ters which have a different AS number than their network. That's how seriously these are different networks.

An ISP is not going to be buying into a hosting center. An ISP is going to get a T-3 or an OC3 into a big carrier's router. So we say, we're congested. They say, oh. Okay, we could in four months give you a port on this router rather than running through a switch. Well, but what router? Oh, the customer attached router, okay put us on the backbone router. Okay, we'll put you on the backbone router. Where's the router? Oh, the backbone router still only has an OC3 into the big carrier network. And still has a hundred megs per second of traffic on it.

They can't provision additional OC3's, because they're not doing a gigabit architecture. The other carriers are all the same way because they are all using Cisco 7500's and a couple of OC3's. But the irony of it is they're charging more, because this guy who wants three megabits per second of connectivity, it's worth it to him to pay 50% more for the bandwidth to not have to worry about running a router. But for us, that's unacceptable.

COOK Report: And the ISP hopes he gets adequate service. Or does he really get adequate service?

Limits to Scalability of the Vertical Peering and Backbone Connectivity Model

Freedman: But he does get adequate service. Because if you have fifty people using 3 megabits per second, most of them don't start using 3 megs per second the next day. Most of them grow very slowly. That's how the carrier makes the model work for the ISPs. But for the big websites huge need is there right now. Therefore, they can't be accommodated in this architecture.

So what does that mean? That means that Akamai was already on the path towards moving out to the edge. Broadband, big ISP's, like that. But some of the, but how do you get into the bigger networks? The answer is not through the hosting centers. The answer is now Akamai needs to start looking like a network — even though it isn't... Now this is where it can hurt your brain.

COOK Report: Yeah, right, go ahead.

Freedman: So we'll get OC12's. Now there are some networks that don't want to put us in their network, but have completely open peering policies.

COOK Report: But you get an OC12, you get this damn OC12 from exactly where to exactly where? You get it from the web farm around to the backbone such that you've got a separate link that may be 100 feet or something?

Freedman: It might be sitting inside Joe's Network, where our rack has more connectivity than the whole data center does, because that's what we need. And the leased line guys, leased lines don't get aggregated as much. And by buying OC12's, you make sure, I guarantee you, if somebody's aggregating an OC12, they're under congested in a lot of their backbone. Most people don't have customer-attached routers you can put OC12's on. Or, if their hosting center is full or they see what we're doing and they don't like it, then we'll rent a room somewhere, put the OC12 in or we'll bring it into another data center. Or whatever, but we will do it in such a way that we begin to aggregate multiple sites, but not grow a long haul network. And I don't want, if I'm sitting at Abovenet, and I have a gigabit to them that I can actually use, I don't want to start bringing OC12's in there, because that means, if Abovenet has a power problem — they never have, but let's say they do — a lot of my connectivity is affected. I don't want to aggregate these things and put them near each other. So it might be inside the same hosting center, it might be somewhere else, whatever, we get an OC12 into their router.

COOK Report: Well, I have one other parallel question. Does the structure that you have provide a means of helping big websites deal with distributed denial of service attacks? Or is that totally orthogonal?

Freedman: Yes, it does. But wait: are you talking networks or customers? Networks or content providers.

COOK Report: Content providers.

Freedman: Oh, yes, sure. Because if someone attacks one of our regions, the region just falls offline. Our network operation center doesn't, isn't required, really. Because the network self-heals.

Now, we don't go out there saying, buy our product and you'll never be down. Buy our product and you're immune. But if someone attacks, let's say they're attacking content provider A and one of our regions in a hosting center is right next to that, that region may go offline. But our algorithms automatically redirects the requests somewhere else. Remember when I told you the mapping, some of the name servers handling each region are actually elsewhere. So if a big hosting center in Bay Area goes offline because of a denial of service attack, it doesn't effect anything in our network and to the extent that the content was pushed onto our network previously, then we still have it.

COOK Report: If back in February, Yahoo was seen from several parts of the U.S., if I remember correctly, to be essentially unavailable. But from other parts in the world, Canada and I think abroad, it was quite available. How come?

Freedman: This is actually a pretty long answer and there's another product that we have that really helps with this, which I'm not sure I can tell you about. But this is a complicated answer to get correct, because Yahoo is a customer and they do buy a few different services from us. But the bottom line is if someone is only serving html from a couple of places in the world, then it has nothing to do with Akamai. Unless they're serving their front page through us, if they get attacked, they're not going to be available, they're going to be offline. But there is a lot that relates to what we do and how we do that provides good protection against denial of service attacks.

COOK Report: So where is all this headed?

Freedman: Well, my group at Akamai is integrated. We have two main functions. We work with R&D, especially in the realms of how the computer is attached to the infrastructure, because that's sort of controlled by R&D, so moving from fast DSL to Gig Ethernet and beyond, that's a network architecture issue. And my group is network architecture. We work with the mapping group to better understand how the Internet fails.

COOK Report: The mapping group being the group that looks at traffic patterns?

Freedman: Yes, there's a group that builds a map that's supposed to be accurate to within seconds of the complete performance reachability information of the Internet. I wish I could spend all my time in that group, but my people are needed to help with the building of the network also, because that's really interesting stuff.

Thus we work with the mapping and performance analysis groups in R&D, but we are, for now, really directing the network architecture and strategy and deployment parts of the company to ensure three things. We want our networks to be burstable, with excellent performance and to also be the most effective at the lowest cost.

COOK Report: But your network's not really a network, it's kind of a virtual overlay.

Freedman: Right. Our virtual overlay network. Our bunches of boxes.

COOK Report: And where you are needing them, you are getting either semi-permanent

or temporary patch solutions within a data center to get from point A to point B through the data center's network or if the network is performing poorly, you're in there to negotiate, how you will go from point A to point B, via detour of this separate OC3 cable if need be.

Freedman: Right, or into the backbone. But we're also deploying peering regions. Because there are bunches of providers out there. And there are some people think, like from your question, who think that we're caching which is OK. What they don't realize is, if they just BGP peer with us, data magically appears. Consequently, we're going to set up a router. We're going to peering points. We're going to be peering with people over peering points and begin looking sort of like a network from that perspective.

COOK Report: So you're extending your virtual network, you're extending it a bunch of different and new ways.

Freedman: Yes. And in Europe and Asia, bandwidth isn't necessarily costly, necessarily, if I do it right, because let's say you're a network that's peered in Asia. And I come to you and say, I can give you ten megabits per second to give to all your peering partners and I'm not going to charge you for it. What's the value proposition? Those peering partners would otherwise have to pay to bring that data in over international circuits. That's very valuable data. I call our boxes sometimes Magic Traffic Generators. I can put them there and make more traffic appear, as long as the quality's good.

So it helps them. When you're in their network, if you can double bandwidth at no cost to you, you all of a sudden look more attractive to peers internationally, because you're giving them much more traffic. Therefore it helps us to widely deploy in peering points there and even if we don't peer with people, just by physically putting our content there, they may decide to give us their good connectivity.

Now in the U.S., let's there's a network A, that doesn't want us to deploy inside their network. But they peer with 50 other networks. They don't want to peer with us because they think they want to compete with us.

But they peer with all these people. Therefore all I have to do is go to the people they peer with. Now let's say you, Gordon, are one of these networks that peers with network A. And I say to you, okay, you're cost is $25 per megabit for these circuits. Let's say I put boxes next to each one of your peering pipes and pay you $100 per megabit. Now you're sending more traffic to that network, so you look more attractive as a peer and you're making some money and you're making three times what your cost is.

That's another way we can deploy. So it's just using creativity to figure out what the best way is to get our customer's traffic to the end users of networks who, in some cases, aren't yet on the page in terms of letting us deploy all over them.

COOK Report: And so obviously they've come to you as someone who, over the last several years, has developed a very, very good knowledge of the players and how it works; who the people are; the technical scalability issues. Plus, if somebody tells you they're going to do such-and-such at a big data center, you know and they know that you know they know, that you can check up on them and hold them to what they've agreed to do.

Freedman: Yes, that about wraps it up.

Some Akamai Traffic Statistics

[**Editor's Note**: An email discussion with Avi Freedman on April 15 - 16 yielded the following postscript..]

Since Akamai bills its content providing web sites in terms of the aggregate content requested from each of the sites world wide, it must keep good track of the aggregate data shipped to its servers at the edges of the Internet. Or, in more traditional e-commerce terms, data delivered to end user "eyeballs."

When Akamai aggregates this data, it does so in terms of the AS (Autonomous System Number) of each network. While there are a few AS numbers that would generally not be thought of in terms of ISPs numbers and a few ISPs that have more than one AS number, ISPs globally tend to have a single AS number for routing purposes. While the size of the Internet can be measured in more than one way - for example by numbers of email accounts - when networks are ranked according to their percentage of Akamai traffic received one can certainly get some idea of their relative size. The top four globally in terms of percentage are UUNET, Sprint, AOL and ATT. The total for ATT here is combined ASNs of ATT Worldnet, Cerfnet, and the old IBM Global Network (now called AGNS). AOL's > 10% is the combination of a number of AS numbers.

Note again that the following networks are listed in terms of where the eyeballs (Akamaized bits are being sent) are, not from where the bits are originating. Only ISPs and other networks that are multi-homed; truly Tier 1; or got an ASN before 1996 or so will appear in this list. Finally one must note that the table counts all traffic that CAN cross any given network. So let's say that on its way to Joe'sNet which terminates 0.20 of Akamai's total traffic, the traffic transits UUNET or Sprint or any other backbone. In such a case the traffic in question is counted both for the transit backbone and for the net where it terminates. The actual totals of the percentages add up to roughly 277% of aggregate traffic meaning that the average IP prefix is multi-homed 2.77 times.

There are 2 networks > 15.00 %,
UUNET and Sprint

There are 2 networks > 10.00 %,
AOL, and ATT

There are 4 networks > 4.00 %,
There are 6 networks > 2.00 %,
There are 4 networks > 1.00 %,
There are 34 networks > .50 %
There are 56 networks > .25 %
There are 22 networks > .20 %
There are 44 networks > .15 %
There are 85 networks > .10 %
There are 61 networks > .075 %
There are 139 networks > .05 %
There are 77 networks > .04 %
There are 95 networks > .03 %
There are 197 networks > .02 %
There are 487 networks > .01 %
There are 196 networks > .0075 %
There are 378 networks > .0050 %
There are 721 networks > .0025 %
There are 885 networks > .001 %
There are 2849 more networks

Traffic to Europe is accounted under each large transit AS. Europe is about 8%, Asia about 3%, South America about .5-1%.

Listed in alphabetical order the 14 networks in the greater than 1 percent but less that 10 percent categories are Abovenet, BBN, Bell Canada, Concentric, Digex, Exodus, Globalcenter, @Home, Level3, PSI, Qwest, Splitrock, Teleglobe, and Verio.

Part Six

ICANN

ICANN's performance in 2000 was just as reprehensible as the year before. Considering it to be a known and potentially dangerous entity about which mere 'humans" can do very little, we repint here only two articles. One is our own analysis from our May 2000 issue. The other is our summary of Michael Froomkin's 168 page Duke University law review article that finds the Commerce Department to have acted illegally in the creation of ICANN.

Executive Summary

Breaking the ICANN Code, pp. 262- 268

Various court decisions are making ever more clear the advantage that possession of the root gives to the US in maintaining its commanding lead in global e-commerce. This is leading to resentment abroad. Given the course on which we are all headed ICANN is likely to be at best a temporary band aide on a festering sore until decisions of foreign courts or governments fracture the US-controlled, authoritative root. We discuss both some of the ways in which this fracture might take place and what impact it would likely have on the Internet's operation.

While a fractured root would certainly not destroy e-commerce, the very fact that it happened would be likely to pop the speculative bubble supporting the stratospheric prices of Internet stocks. It would demonstrate that a globally-unified forward march of the global economy running on internet "rails" is only a pipe dream. Many investors and VCs would be forced to rethink the price value equations on which their actions have been based. Should contention over the root get serious enough to throw prices of Internet stocks into a nose dive, the United States would loose far more than any other nation in the world. This is very likely what John Patrick and Vint Cerf and Esther Dyson had in mind when they asked the venture capital community to contribute to ICANN last summer cryptically warning that if ICANN failed e-commerce would also fail. Certainly the ongoing uncertainty of how much of a global market for business to business e-commerce would be easily reachable in the event of trouble for the authoritative root would take the buzz off of most e-commerce business plans.

We arrived at the above conclusions after pondering Ed Gerck's essay "Thinking" (April *COOK Report,* pp. 23-25). We find Gerck's article to be a useful point of view for analyzing some unresolved issues relating to ICANN and the Department of Commerce on the one hand and the DNS and alleged need for a single authoritative root on the other. Gerck sees DNS as the major centralized component of an otherwise decentralized Internet. In his essay "Thinking" he says that some of the choices made long ago in the design of the DNS not only make it depend on a single root but also "without the DNS there is no email service, search engines do not work, and webpage links fail." DNS is " the single handle of information control in the Internet. And, in the reverse argument," it is "its single point of failure."

With something as powerful as the Internet, everyone wants more and more to seize control, if only to keep others from controlling it. It certainly can be argued, as Gerck does, that the struggle for control of DNS has become the focal point over the last four years of a diverse coalition of actors (trade mark interests, IBM, ATT and others) that have gathered together to form ICANN. It is generally assumed under US law that the organization which controls an entity bears legal responsibility (liability) for the use of its power. Gerck suggests that under the conditions of a single handle of control over the Internet, the controlling organization's liability is potentially total. Thus given the nature of ICANN's use of DNS as a means of grabbing control over the Internet, the liability facing ICANN and anyone else who would emulate it is essentially unlimited. As a result, in structuring ICANN it has been necessary to insulate all players from the consequences of their otherwise unlimited liability.

We have taken Gerck's essay and used it as a template on which we have applied our own knowledge of ICANN. This process has helped to bring a number of issues into focus for the first time. In its eagerness for control those who have promoted ICANN have taken all the critical administrative infrastructure of the Internet DNS, IP numbers and protocols and dumped them into the single ICANN basket. ICANN

But having all our eggs in one basket and having in the DNS a single point of failure creates the kind of prize that, as long as we still have national economies competing against each other, the US government and its major corporate allies will do what ever is necessary to protect from foreign capture, or even from foreign influence. Since ICANN is the basket holding all the eggs it, in the meantime, must be protected from its unlimited liability by being made virtually unsuable.

But in order to make ICANN unsuable, its backers have had to create for it an arbitrary structure that renders it immune from the inputs of those communities that it is supposed to serve. This arbitrary structure has in turn prevented ICANN from inheriting the political legitimacy within the Internet community that Jon Postel's exercise of these functions once enjoyed. ICANN follows a carefully scripted routine that supports its role as guardian of all the Internet's administrative eggs that have been in its single "basket." This scripting greatly angers those who having mistaken the ICANN process for being one of actual openness have invested their time in hope of influencing the outcome. However the play acting also serves ICANN interests in that it can be spun by ICANN's public relations firm in such a way that the casual press lacking the time and ability to do its own research may be fooled. Therefore ICANN has bought the administration some short term time to regroup and maneuver.

What we have done in this article is demonstrate (1) why ICANN can be nothing more than a temporary fix (2) how ICANN is likely to fail (3) why the consequences of this failure will hurt the United States more than it will hurt other nations, (4) why from ICANN's efforts designed at all costs to shore up what is really an untenable effort to maintain long term central control over Internet addressing there needs to be a switch to efforts aimed at placing in the hands of each user the means by which he or she shall be able to address and find Internet objects.

ICANN was created as a diversion on the part of Ira Magaziner who conveniently left the administration and returned to private consulting as soon as it was established. It is a smokescreen cleverly designed to give the illusion to the rest of the world that the US is transferring control of administrative functions over the net to a world body where the Europeans and Asians would be led into thinking they could play a significant role in policy making.

And indeed just so long as they don't try to grab the root, American policy is to play along with the Europeans and Asians and acting through ICANN do such things as granting them direct control of their own country codes, and the power to enable their corporations to have preferential treatment over domain names on the excuse that such names can be treated as trademarks. Many other powerful groups have been given an opportunity to play in the great ICANN charade.

As long as ICANN is there, it gives the impression that others besides the US government will be allowed a role in root server policy making and control. In reality the continued heavy hand behavior of Roberts and Dyson has made it possible to drag out the ICANN foundation process for another year getting it conveniently past the upcoming touchy US Presidential elections. As a result the Clinton Administration has been able to extend the dual relationship of the ICANN DoC cooperative agreement.

The extension makes it possible to preserve ICANN as a maneuver designed to deflect attention from the stark fact that without ICANN, the US administration would seize the root servers by force rather than loose control. This is the secret of why ICANN cannot be allowed to fail. ICANN's central

purpose is to divert attention from the fact that the Clinton administration has made a decision to treat the root servers as a strategic telecommunications resource over which it is perhaps even prepared to use the police power of the state to protect from falling into the wrong hands.

It would be encouraging to see some interest in Washington in the incubation of the understanding necessary for the internet and e-commerce to cooperate in working its way out of the win-lose control situation in which it finds itself. The route of control has been tried. As we have shown in this discussion, not only has it not worked, it also looks to be untenable on a long term global basis. It is to be hoped that if our policy makers understand that we are likely to loose more than anyone else in a struggle to maintain our control, they may also come to understand that they have the most to gain by removing all possible levers of control from everyone's grasp. If it becomes clear that no single entity can hope to control the Internet, many strains in the present system could be quickly dissipated. We are a "can do" nation. If the administration were to understand that everyone would have more to gain from such an outcome we believe that there is adequate talent available to ensure a success.

ICANN Deemed Illegal, pp. 17 - 21

We summarize and comment on Michael Froomkin's 166 page 711 footnote long landmark paper: "Wrong Turn in Cyperspace: Using ICANN to Route Around the APA and the Constitution" to be published by the *Duke University Law Journal*, October 2000, Volume 50, No. 1. Froomkin's indictment in his opening paragraph is succinct: *"The United States government is managing a critical portion of the Internet's infrastructure in violation of the Administrative Procedures Act (APA) and the Constitution. For almost two years the Internet Corporation for Assigned Names and Numbers (ICANN) has been making domain name policy under contract with the Department of Commerce (DoC). ICANN is formally a private non-profit California corporation created, in response to a summoning by U.S. government officials, to take regulatory actions that the Department of Commerce was unable or unwilling to take directly. If the U.S. government is laundering its policy making through ICANN, it violates the APA; if ICANN is in fact independent, then the federal government's decision to have ICANN manage a public resource of such importance, and to allow - indeed, require - it to enforce regulatory conditions on users of that resource, violates the nondelegation doctrine of the U.S. Constitution. In either case, the relationship violates basic norms of due process and public policy designed to ensure that federal power is exercised responsibly.*"

We believe that it is very important to use Froomkin's compelling insights to educate both citizens and the executive and legislative branches of the US government. We need to understand quickly what has happened and why we should "be afraid." Out of such education it is to be hoped that legal or legislative redress may be found. In a brief interview with us Larry Lessig explains that Froomkin has provided a road map for legal that under the right circumstances could be used to compel ICANN to change its ways. The almost final draft of the Froomkin paper is a 1.1 meg pdf file at http://www.law.miami.edu/~froomkin/articles/icann1.pdf We also offer comments on the recently concluded ICANN membership at large elections. We suggest that the determination of ICANN's fate is a decision that will impact whether we live in a nation state with a legal system accountable to individual citizens or a corporatist state that is run on behalf of the private profits of global corporations. If it is the latter we will find that we have given up the critical philosophical foundation on which our nation was built. That our government exists first and foremost to preserve and protect the rights of each of its citizens.

Cracking the Code: an Analysis of US Internet Governance, E-Commerce and DNS Policy

Why US Dominance of E-commerce Indeed is Dead if ICANN Fails & Why the US Has Most to Lose from Continuing a Policy Founded on Indefinite Control of the Root

We have found Ed Gerck's essay "Thinking" (April *COOK Report*, pp. 23-25), to be a useful point of view for analyzing some unresolved issues relating to ICANN and the Department of Commerce on the one hand and the DNS and alleged need for a single authoritative root on the other. Gerck sees DNS as the major centralized component of an otherwise decentralized Internet. In his essay "Thinking" he says that some of the choices made long ago in the design of the DNS not only make it depend on a single root but also "without the DNS there is no email service, search engines do not work, and webpage links fail." DNS is " the single handle of information control in the Internet. And, in the reverse argument," it is "its single point of failure."

With something as powerful as the Internet, everyone wants more and more to seize control, if only to keep others from controlling it. It certainly can be argued, as Gerck does, that the struggle for control of DNS has become the focal point over the last four years of a diverse coalition of actors (trade mark interests, IBM, ATT and others) that have gathered together to form ICANN. It is generally assumed under US law that the organization which controls an entity bears legal responsibility (liability) for the use of its power. Gerck suggests that under the conditions of a single handle of control over the Internet, the controlling organization's liability is potentially total. Thus given the nature of ICANN's use of DNS as a means of grabbing control over the Internet, the liability facing ICANN and anyone else who would emulate it is essentially unlimited. As a result, in structuring ICANN it has been necessary to insulate all players from the consequences of their otherwise unlimited liability.

As ICANN has become the sole authority for the administration of DNS, it has acted to separate itself from the consequences of the administration of that authority by franchising it out to one registry and 100 registrars. The entities act as ICANN's agents and are give the ability by means of their franchise agreement with ICANN to act on ICANN's behalf. With the insertion of three new parties (Registrars, ICANN, and US Government into a process that last year had only two parties (registrant and registry) a multi-party process has been established. Now instead of one entity (registrar) having liability for what it does to the interests of the registrant, ICANN is able, for the time being at least, to escape the potentially unlimited liability. The registrar and registry can say they are just doing what ICANN forces them to do. In turn ICANN can say that it cannot supervise every step made by over 100 registrars and one huge registry. Furthermore in the case of relations between Network Solutions Registry and ICANN each entity can skillfully sidestep the necessity to deal directly with each other since the Department of Commerce has inserted itself squarely in the middle

The Clinton Administration has quite skillfully freed itself from the controversial position of overtly controlling or otherwise regulating the internet by setting ICANN up as a private industry organization designed to achieve perfected and total regulatory control over the internet through putting IP numbers, and protocols in the same single regulatory egg basket as DNS. The liability however is too hot for anyone to handle. In the case of ICANN the very government that set this up has solved the hot liability problem by approving the design of a set of bylaws and procedures that makes ICANN, as an organization, accountable to no one. To the great consternation of many people ICANN simply does what ever it chooses. Of course as Gerck also suggests it is in the interests of the US Government to see that control of DNS does not fall into the hands of some foreign power or powers whose interests are inimical to those of the US. ICANN under the control of the Department of Commerce places ICANN on a leash with the illusion but not the reality that ICANN will soon be set fully free.

Therefore, in this article we contend that this total divorce of ICANN from the bottom up, consensus-based behavior that it was allegedly founded on is the result of, as Gerck suggests, as struggle for the single handle of control over the DNS and over the single control point for DNS - the single authoritative root for the generic top level domain names - especially com, net and org.

We observe that our own government was founded on a separation and balance of powers, in other words on the assumption that the American political system would rest on diversity and not on a single point of control. However, given the perhaps inadvertent single handle of control for the Internet, the Clinton administration has taken the easy way out and for the protection of national security interests has allowed the administration of all technical aspects of the internet to fall into the hands of a single unaccountable entity whose only interest is in control achievable through the manipulation of DNS and the root.

We also contend that the enormous power involved in control of the root acts like a magnet to draw the attention of all those who are dissatisfied by current policies or by the fact that their hands are not the ones on the levers of control. Unfortunately instead of trying to ascertain whether there may be technical changes that could render the Internet controllable by no one, as Gerck suggests, the administration has placed all its (and our) eggs in the ICANN basket.

We believe that the decision to vest the location of control of the DNS in ICANN has required a chain of supporting decisions to render its control viable over time. However this also has meant that in order to gain immediate viability it has been necessary to create an ICANN whose unlimited liability is protected by isolating it from any possibility of influence by those whom it is supposed to serve. Such an ICANN has to pretend to be what it is not and without any genuine support (save for that of its control seeking founders) is unlikely to survive over time.

Let's restate what we just said. The point we are making is critical to breaking the code and understanding what the ICANN process is really about. All eggs in one basket and a single point of failure creates in DNS the kind of prize in control of the single authoritative root that, as long as we still have national economies competing against each other, the US government will do what ever is necessary to protect from foreign capture, or even from foreign influence. ICANN itself meantime must be protected from its unlimited liability by being made virtually unsuable.

But in order to make ICANN unsuable, its backers have had to create for it an arbitrary structure that renders it immune from the inputs of those communities that it is supposed to serve. This arbitrary structure has in turn prevented ICANN from inheriting the political legitimacy within the Internet community that Jon Postel's exercise of these functions once enjoyed. ICANN follows a carefully scripted routine that supports its role as guardian of all the Internets administrative eggs that have been in its

single "basket." This scripting greatly angers those who having mistaken the ICANN process for being one of actual openness have invested their time in hope of influencing the outcome. However the play acting also serves ICANN interests in that it can be spun by ICANN's public relations firm in such a way that the casual press lacking the time and ability to do its own research may be fooled. So ICANN has bought the administration some short term time to regroup and maneuver.

What we intend to do in this article is demonstrate (1) why ICANN can be nothing more than a temporary fix (2) how ICANN is likely to fail (3) why the consequences of this failure will hurt the United States more than it will hurt other nations, (4) why from ICANN's efforts designed at all costs to shore up what is really an untenable effort to maintain long term central control over Internet addressing there needs to be a switch to efforts aimed at placing in the hands of each user the means by which he or she shall be able to address and find Internet objects.

The Weakness of the Current Game Plan

The administration is not getting good guidance because it is asking the wrong questions. In our most charitable interpretation, prior to the failure of the WTO in Seattle in December 1999, the administration may have thought that, with ICANN, it had a single body which it could use to run the Internet in the name of a single global economy. We contend however that in an election year with the stock markets at record heights and no matter what its actual intention, the administration is playing a dangerous game by turning loose an entity designed to run the Internet in the interests of a small number of global corporations and at the expense of vast numbers of small businesses both in the US and abroad.

The world is more and more a collection of nation states with different languages, cultures, religions and economies despite all the attempts of media pundits to talk about globalism, global culture and a single global economy - just ask the citizens of Seattle. In this context with the election of a US president and congress only six months away, the administration is not about to give away the control of the DNS root - not even to a global body ICANN that it put together in such a way that it would be almost impossible for foreign interests to control and use against the United States.

Let us state outright what we have effectively implied above. One way to keep the enormous liability given by control of the DNS from landing squarely on the shoulders of ICANN or the US Department of Commerce is to do precisely what is being done now. Keep control of the root while all the time grooming ICANN as an organization that "real soon now" will achieve globalization of internet control by giving away to foreigners the United States' most important cyberspace strategic asset, namely, control of the DNS root. We have just seen how this pattern played out in ICANN's March 7th to 10th Cairo meeting with the statements that four members of the current board would stay on another year beyond their already expired terms and the assertion that, in view of unfinished business, ICANN would ask the DoC to extend its cooperative agreement beyond the September 30th, 2000 expiration date.

ICANN currently has all the goods needed to become a global Internet regulatory body save one, control of the root. In the mean time it has set in motion the necessary machinery for allowing trademark interests to run the Internet's naming system. ICANN's world girdling road show has been played against a back drop of the most arcane and complicated politics in a way such that observers have only two choices.

Those who are part time specialists can accept the official version of events and be reassured that ICANN is doing nothing more serious than administering otherwise obscure parts of the internet's plumbing. The cacauphony of complaints and accusations by insiders surrounding ICANNs twists and turns, must seem dreadfully confusing to general observers unable to give ICANN matters much attention. To those of us who have become nearly full time captive observers of how ICANN has manuevered the process, it seems that the current confusion may be of value to the game that is being played. Namely one of encouraging European and Asian interests to hope that the US will give up control of the root when politically there is no way that this can happen under the current DNS technical system. Given the current confusion there is no way of telling under which shell, ICANN or DoC's the winning cards in the on going game of three card Monte lie. Since as we have seen with its continual changes of its bylaws, ICANN's rules are to have no rules, it is all a very clever game of deception. So far the rest of the world has played along with the charade.

Why Can't ICANN Be Allowed to Fail?

At various times John Patrick, Vint Cerf and Dave Farber have all claimed that ICANN must not be allowed to fail or the Internet will be in danger and electronic commerce will collapse. We have repeatedly asked them to explain their veiled warnings. They have repeatedly refused. We believe that we have figured out why: ICANN was created as a diversion on the part of Ira Magaziner who conveniently left the administration and returned to private consulting as soon as it was established. It is a smokescreen cleverly designed to give the illusion to the rest of the world that the US is transferring control of administrative functions over the net to a world body where the Europeans and Asians would be led into thinking they could play a significant role in policy making.

And indeed just so long as they don't try to grab the root, American policy is to play along with the Europeans and Asians and acting through ICANN do such things as granting them direct control of their own country codes, and the power to enable their corporations to have preferential treatment over domain names on the excuse that such names can be treated as trademarks. Many other powerful groups have been given an opportunity to play in the great ICANN charade. And the cast of characters who have fallen for the deception is quite impressive.

Some like Dave Farber (who should know the game being played) go along with the dance. For example on Thursday March 9th Farber complained to his IP mail list that Jeri Clausing of the New York Times should not continue to call ICANN an oversight body. She should let Times readers decide for themselves. What is bizarre about this episode it that several months ago Farber sent out without comment a similar message. For his effort, Farber was chided by Esther Dyson. Now with ICANN's oversight role clearer to everyone than ever before, Dave Farber takes issue with Clausing's obvious description. Why will this intelligent man go to such lengths to protect and shield ICANN from the facts obvious to all disinterested onlookers?

Because, we suggest, that as long as ICANN is there, it gives the impression that others besides the US government will be allowed a role in root server policy making and control. In reality the continued heavy hand behavior of Roberts and Dyson has made it possible to drag out the ICANN foundation process for another year getting it conveniently past the upcoming touchy US Presidential elections. As a result the Clinton Administration has been able to extend the dual relationship of the ICANN DoC cooperative agreement.

The extension makes it possible to preserve ICANN as a maneuver designed to deflect attention from the stark fact that without ICANN, the US administration would seize the root servers by force rather than loose control. This is the secret of why ICANN cannot be allowed to fail. ICANN's central purpose is to divert attention from the fact that the Clinton administration has made a decision to treat the root servers as a strategic telecommunications resource over which

it is perhaps even prepared to use the police power of the state to protect from falling into the wrong hands. With ICANN in existence those who want control of the root have two obstacles to deal with. On the one hand the goals of ICANN's backers and on the other what the US government might do in the future to maintain US control of the root.

Some History

It was not always this way. Through the efforts of IAHC, control of the root was almost lost in 1997 before Clinton, Gore, Magaziner wised up and the turf on which the game was being played changed. To stabilize the situation Ira ran around the globe convincing folk that the US would share the "cookies" with them and that "NewCo" soon to become ICANN was the vehicle to accomplish the sharing. Relatively few people noticed that in the last months of Jon Postel's life Ira made it very plain to Jon that the root was regarded as a strategic economic and security resource that the feds and not Jon controlled.

ISOC wasn't pleased to hear this assertion of federal authority, so in the Spring of 1998, to placate ISOC, it was given an opportunity to run the ICANN strategy. It could be argued that during the next 18 months ISOC tried to grab the single handle of control and alienated everyone except for the ISOC faithful. The controversy, as ICANN was born (in still undefined circumstances) and began to do its business was fine with the White House. It would give the feds the opportunity to extend DoC control and provide both the opportunity and the excuse to build a DoC claim to the root. Finally it would permit the extension of its indirect control over the root indefinitely while it waited " just a few more months" for ICANN to settle down and prove itself worthy of accepting the root. But as long as ICANN was controversial, there was no need to transfer the root. Therefore as a root-holder-in-waiting ICANN takes the pressure off what otherwise might be seen by the rest of world as a naked assertion of the strategic interests and power of the US.

ICANN was needed to divert attention from what was really going on. The White House couldn't afford to loose ICANN in the summer of 1999 and have attention shifted totally to it and the Department of Commerce. That is very likely why when in June of 99 Mike Roberts wrote Tom Kalil at the White House and asked Tom to help raise money to keep the diversionary deception going. Tom welcomed the ICANN foot soldiers to his office and pledged his help. The continued existence of ICANN is of great value to the US government because it enables it to avoid being the sole focal point for the liability involved in being responsible for the sole point of failure of the Internet.

Fast Forward

But suppose ICANN fails. A collapse of ICANN would end the game of "root root who's got the root" because Commerce then stands as the lone guardian of the root. With the ICANN deception ended, we suggest that the root will splinter because the rest of the world will not accept, without hope of change, a global system having a single point of failure and control under the sole sovereign power of the US government. The American desire to maintain a system designed to serve the needs of a tiny elite of big government and big corporations at the expense of foreign economies and ordinary citizens will collapse for the same reasons that the WTO went down in flames in Seattle in December.

The number of scenarios by which this could happen are of course vast. For example: No! China says we will run our own root, France says the same thing. Others follow. Now a splintered root will not mean the end of the Internet. Email and web sites will still work, but depending on how the root splintered and for how long, they may work for far fewer people. E-commerce will be less global and more local. The United States then might find that the global playing field which it presently dominates has been leveled. But trillions of dollars are at stake because already absurdly over valued Internet stocks could suddenly find themselves presiding over a diminished market for their goods and services. Market capitalizations of many companies could melt down if the vision of a single global market for e-commerce were shaken.

Consequently the stability of the global economy is endangered by this bizarre struggle for control over the power inherent in the Internet's single point of failure. This reminds one of the Ring of the Nibelungen. In Wagner's ring cycle - everyone kills nearly everyone else because each must possess the power of the ring (err.... root). No one has the creativity to ask whether removing all possibility of control by a single entity could work to the global advantage of all stakeholders? (Anyone for tossing the ring back into the Rhine as happens in Die Gotterdammerung ....Twilight of the [net] Gods?) In the meantime we still contend that ICANN is needed to dupe some folk into believing that the game is about something other than "root, root who's got the root?"

ICANN is running out of steam. All the keystone cops vignettes of Esther Dyson's valiant efforts to explain why ICANN is run like a banana republic where any actions that smack of democratic respect for due process are forbidden, are wearing thin. ICANN is designed, on the one hand, with a schzoid charter to enable American security and big business interests to hold on to the root no mater what. While at the same time, the public's attention is diverted from the corporate-governmental real politik agenda of naked self-interest by the false assertion that ICANN is really about American democracy and bottom up processes. Such claims are dogs that won't hunt. But sanity will not happen because Commerce can't afford to have ICANN succeed. It cannot afford it because Commerce cannot afford to have to give ICANN the root on a permanent basis. The arrogance of ICANN's pronouncements will therefore continue. And ICANN operating in defense of its control mission above all else will get the respect of those who are ready to exert control at all costs. The rest of us will eventually figure out that we are being manipulated.

The question then becomes why hasn't some one moved long ago to eliminate such a single handle of control? One answer is engineering inertia. It still works so don't fix it. But perhaps the other answer is that a significant part of the federal mission is control, the administration people may be loath to even think about whether there could be a solution that would bar everyone from control and whether such a solution would have a more beneficial impact on the world economy than their pursuit of control through ICANN.

Part Two

Some Thoughts on Possible Fates of the Root

While we have of course not been able to make anything like an authoritative survey we have made several inquiries among members of the Internet technical community as to how stresses and strains on the root may play out over the next year. Concern over the fragility of the authoritative root servers is a widely shared but seldom if ever publicly acknowledged. It is felt however that the fragility is largely legal and, or political.

Some of the more conservative engineers speak of changes to the roots as an incipient Armagedon. Bind 9 is expected to be released by year's end. It will allow more and more root server data to be cached locally, so that the roots have to be queried as seldom as possible. Even so, to these folk the single authoritative root is so fragile that we dare not talk about its problems to anyone. Furthermore to enable the continuance of a

single control point they compound the problem of eggs in one basket by decreeing for example in mapping telephone numbers to objects in the internet, nothing other that DNS will be used because they can't admit any possibility that anything else that could used for locating something on the Internet. To these folk the single authoritative root is the sacred cow of all sacred cows in the Internet. But legal or political action could reveal its immutability as myth. To some people the problem is not DNS but creating a situation where the 13 servers are no longer authoritative and where one or more of the 13 may decided to point to a regional non authoritative root. People ask how can such a global resolution system be protected against political interference?

The point is that it very likely cannot be protected. Since the DNS infrastructure is here the breach would likely begin within the US. Say a US agency goes to court against an off shore banking entity that it doesn't like or a gambling operation, or a prescription drug selling operation and the plaintiff gets wise. If the name is .com, .org, or .net, the US Agency will ask for removal of the offending name from Network Solutions Registry Database.

Suppose however that the rogue or "outlaw" operation is using one of the roughly 240 country code TLDs, the court could ask the administrator of the country code data base to remove it from the country code data base. If the administrator of those country codes was a business located on American soil, he or she would likely comply with the orders of the American court. If that administrator were located on foreign soil and were hostile to the position of the American court and refused then the only remaining viable option for the American court is to order network solutions to remove the country TLD from the root. If it were a small country code like the Pitcairn Islands the argument would be made that it never would be missed. If it were some what larger (Niue or Tonga say) with hundreds or thousands of businesses registered there, then all hell would break loose as those organizations would take legal actions of their own to defend their identity in cyberspace. NSI might find itself stuck in a withering crossfire.

Since the root sits in the US as a national strategic asset, NSI, after consulting with its DoC and ICANN masters, would very likely have to obey an American court order to remove a dot com registrant from its database. Removal of an entire country code fromthe root would be less likely and would certainly involve NSI' s asking DoC for permission before it did so.. A DoC backed order to remove a single country code second level domain registration from the country code zone file would be more like likely. In any case it is possible to imagine a future action from within the U.S. legal system that brings about the death penalty for an "outlaw" operator that happens to be a tld country code..

But we must ask what would NSI and DoC do if asked by a court in Britain to remove a second level domain name from the ICANN registry? Let's make things worse. What if a British company wants the name of a French company pulled and the French company wants it left in and gets a French court to tell NSI to leave it in? What does NSI do then? As foreign governments come to understand that ICANN is primarily a shell game designed to mask the dominance that possession of the root gives the US, how long before each country asserts its national sovereignty by setting up its own national root?

Should such a scenario unfold it would fracture the authoritative root and set back the Internet as a mechanism for global e-commerce for some number of years. Angry politicians in Britain and France and a host of other countries might like that just fine because it would harm the current American dominance of global e-commerce far more than it would harm their own economies. Also they could argue that such action would encourage the building of Internet commerce which is more local and regional in character, thus positively assisting those foreign nations whose local businesses are being stifled or diverted by American internet dominance, not to mention the suppression and diversion of their own commercial legal and tax and regulatory systems.

Examples of foreign discontent are popping up all over. Consider Declan MCCullagh's March 14 *Wired* news report: - http://www.wired.com/news/politics/0,1283,34937,00.html

UK Wants Tighter E-Trading Laws "It's not that Phillip Thorpe hates the Internet, not exactly. But he sure wishes the problems it causes him would vanish as quickly as Bill Bradley's presidential ambitions. Thorpe has one of the most difficult jobs anyone can hold nowadays: He's a regulator. And not just any regulator, but the head of the British agency responsible for overseeing the financial services and securities industries. In the U.K., London-based Financial Services Authority can shut down any unacceptable Web site that doesn't follow British law, but the agency's reach is limited."

"Thorpe is concerned that British subjects can connect to any Web site anywhere in the world to deposit money or spend their cash on risky — though perhaps profitable, and probably unregulated -– ventures. There's precious little Thorpe or his FSA cohorts can do about it. "Someone 10,000 miles away simply refuses to recognize your authority," Thorpe griped during a speech Monday afternoon at the Global Internet Summit at George Mason University. "It's set to get worse if nothing changes." . . . . "The problem is made worse by national governments having very difficult ideas [about what should be legal]," Thorpe said."

The rest of the news report concluded that Thorpe agonized about what might be done to salvage the situation. His ideas focused mainly upon changes in the software structure of the web including modifications of http to include a separate family of regulated financial URLs. We believe that if he pursued this direction of thinking that he would soon see that such attempts to control are unworkable.

Regulatory frustration will be one source of discontent with American hegemony exercised over the single handle of control. Without far better leadership there is another source of discontent that will increase pressure abroad for political challenges that are designed to right the imbalance likely directed at the root. This source is the geographic advantage that now may be given to dot com registrants located in the US.

Consider also the March 15 press release from one of the UK law firms involved in a case decided by a Virginia court on March 5, 2000. "US SEIZES JURISDICTION OVER DOT.COM COMPANIES · US companies able to seize worldwide registrations · US trademark owners able to have domain names of others expropriated,"

"In its Decision of Friday 5th March 2000, Caesars World, Inc -v- Caesars-Palace .Com and others (Civil Action No 99-550-A), Judge Albert Bryan effectively decreed that his court would be the arbiter of the property rights in respect of all the approx. 7,000,000 .COM, .EDU, and .ORG top level domains registered. Although the Decision dealt solely with the motions of two of the defendants to dismiss an action for lack of jurisdiction, the effect of the Decision is to open the floodgate to litigation by holders of US trademarks against domain name proprietors based outside the US (or at least outside Virginia. . . . . However, for the bulk of the Internet community the issue is not one of the substantive issues but one of jurisdiction and being subject to the jurisdiction of the US courts."

David Flint, Partner in MacRoberts Solicitors Glasgow (who advised one of the defendants in this case) comments: - "This decision may have the greatest effect ever of any US cyberlaw case to date. Effectively all those international businesses using Top Level Domains (which is now the norm) may find their domains being challenged by US businesses. Even if ultimately successful, the cost to business in fighting challenges before the US courts is not insubstantial and the possibility of little or no recovery of ex-

penses may lead to a swing in the balance of power to the powerful US trademark lobby. Developments in this case will be watched with interest by world-wide Internet users".

Might the DNS Root Split?

As people outside the United States are increasingly realizing, the Internet e-commerce playing field has a distinctly American tilt to it, the ICANN US Department of Commerce shell game notwithstanding. On March 14, we initiated, on the cybertelecoml list, a discussion of some of the points of uncertainty facing the global administration of DNS.

COOK Report: What is to prevent China from setting up its own root and legally requiring all Chinese ISPs to carry its root.cache? Suppose China were to grab a copy of the NSOL/DOC/ICANN root zone files and remove offensive domain names from it before putting it into its new Chinese government approved root,

David Conrad former director of the Asia Pacific IP number registry and currently working with Paul Vixie in an effort to commercialize the BIND software on which DNS servers throughout the Internet depend, responded to our question: Nothing at all, in fact... In a country like China, it is also (theoretically) possible to force folks to point at a certain root. Of course, this would merely slow people down if they really wanted to get to the server represented by the name.

Fred Goldstein (internet analyst formerly with BBN): "Exactly. China presumably has strict regulation of its ISPs, unlike western democracies. So they could muck with DNS as part of the firewall function, and, in theory, block users from accessing DNS ports outside of the country. Of course a determined user could arrange for zone files to be, say, mailed in, but that's the usual cat-and-mouse game that people play with censors. Or users could access desired sites by number, or via a proxy, until the proxy is blocked, etc."

"But in the USA and other countries that don't regulate the Internet per se, the only reason that ICANN or NSOL has any authority at all is because people accept it. It's a free market, so if for some reason an ISP wanted to point its own DNS at, oh, Alternic or something else, then it could, unless its user contracts specifically prohibited it (which of course said ISP wouldn't do). And for that matter, and end user can set his own DNS pointer at anything willing to serve him, not just his own ISP's. So if there were an alternate root, it would have the authority granted to it by the number of people pointing at it (directly or indirectly). Sort of like a prophet standing on a soapbox — he can say whatever he wants; his authority is based on how many people take him seriously."

"In practice, I think it all pretty much comes down to whatever Bernie Ebbers wants, or maybe Bernie, Steve Case, and a couple of other big ISPs. They point their DNS at the Official ICANN-sanctioned Root, and it counts. They point elsewhere, and ICANN's eggshells are empty. Not a likely scenario, but not one dictated by the gummint. After all, there's no law directing them to follow anybody's DNS. (For a govt. contract, they might have to, but they can always give those customers a different DNS to point at.)"

COOK Report: Suppose France did the same and added French language TLDs of its choosing to its root.

Conrad: One of the issues that is likely to face ICANN in the very near future is the fact that several countries, particularly in the AP region, want to use non-US-ASCII for domain names (e.g., daigaku.nihon (== ac.jp) encoded in Japanese). If ICANN does not address this issue, I suspect you'll see country based root nameservers established in fairly short order...

COOK Report: Might events like these cause a look at some technology issues surrounding the DNS?

Conrad: There are numerous technology issues surrounding the DNS, however I don't expect this is the right mailing list to discuss them.

COOK Report: Why would a foreign government take the extreme step of lessening its connectivity globally? We suggest that it might do so in order to gain more control over what was happening on its own soil. After all, despite the vaunted global economy, nation states STILL exist.

Conrad: This would be unnecessary. The fact that a country can have their own language/character set names inserted into the DNS would likely create a strong market incentive to do just that — people who don't use English normally would prefer to not have to remember the English words for the places they want to visit on the web. Internet doomed, MPEG at 11:00. When I first started mucking about with the Internet (ARPANet in those days), I was absolutely stunned that it worked at all. Now that vast sums of money depend on the Internet, I suspect the owners of those vast sums will do whatever is necessary to keep the Internet from breaking.

COOK Report: at a high level I hear you saying that if you muck around with the current authoritative root and break it there are lots of way to get around such a break.

Conrad: As Fred noted, the DNS (as well as the address) infrastructure works 'cause everybody agrees that it works — Mitch Kapor called it "the Tinkerbelle effect". Anyone can do anything they want, however if you want a consistent and universal effect, everybody has to do the same thing. As most people want consistency, most people do the same thing.

COOK Report: some tend to see fracturing of the current root as a major disaster for the net and for e-commerce, but, when one tries to find out why they are so alarmed they tend not to respond.

Conrad: If you have inconsistent views of the DNS, one person's www.foo.com results in a different Internet site that another person's www.foo.com with all that implies. This would be a disaster if the multiple roots are not coordinated in some fashion (that is, moving the root into an administrative arena instead of the DNS protocol arena).

Note that this isn't really an issue for the internationalized DNS folks — by definition they won't be any issues of coordination, the character sets are entirely disjoint. Of course, if IBM wants IBM encoded in Japanese, they get to talk to/pay the folks who are administering the Japanese character set TLD.

COOK Report: I am suggesting that if a US court tells NSI to remove the second level domain from of an American company .... say "hotporn.com' from its database and hence from the root....NSI might well have to comply. But if a British or French court raised the issue it is less likely that NSI would have to do anything.

Conrad: Right. As a decision by a US court would have little affect in .UK or .FR.

COOK Report: Also I was not suggesting any kind of trans national retaliation....merely foreign frustration if because the root is in the US, American based plaintiffs could get their way when foreign ones could not.

Conrad: The root is, fortunately, fairly boring. It has the gTLDs and the ccTLDs. As such, it is unlikely to be a target of frustration (at least in this context).

COOK Report: Fred and David: let me see if I can bear right down to the nub of the issue. Are you one or both saying that you don't see the presence of the root authority in the hands of the US government as a situation where any court or legislative action (eg. Congress critters pass a law that all off shore gambling sites "dot com names" get removed from the NSI database for example) might cause the kind of strain where foreign the politicians and their government would be willing to sacrifice global consistency for

more control within their national borders. The result being a decision to set up a competing root within their borders?

Conrad: As NSI is the administrator for the current set of gTLDs, I'd argue the organizations that register in those TLDs are implicitly agreeing to be bound by US laws. If they do not agree, they can register in a non-US-administered TLD (e.g., .TO) which may have different laws and/or regulations.

Can action by the US Gov't cause political strain? Of course. However, I'd argue that non-US Gov'ts have already implicitly or explicitly acknowledged the US's authoritative role in the administration of the DNS by not objecting to Ira Magaziner unilaterally declaring the US had this right and accepting and/or participating in ICANN. I'm doubtful that the US Gov't would do anything stupid enough with the DNS to rattle the foreign gov'ts (perhaps I'm too optimistic?). Will the role of the US gov't in the management of the root result in pressure for the creation of alternate roots? Maybe, but I'm doubtful. More likely to result in alternate roots is either the actions (e.g., coming up with offensive new TLDs) or inactions (e.g., not dealing with the question of internationalizing the DNS) of ICANN.

COOK Report: What about the .coms which are located out side the US? What is the percentage? Is it 20 %, or 30%, or even more than 35%? If these dot coms start getting screwed because they are not physically in the US, is there likely to be any recourse for them? Or can the US government possess the root indefinitely without creating a foreign backlash, especially if the ICANN dispute process works predominantly in favor of the US based holders of US trademarks?

Conrad: .COM/.NET/.ORG are US-based domains. Registrants in those domains must deal with the implications of that. [Also note that] The root is relatively static and (should be) fundamentally boring. I have trouble seeing how US trademark law is applicable to names in the root. There are about 250 of them, they are 2 (~240), 3 (7), and 4 (1) characters in length and don't run much risk of confusing consumers (I would imagine).

As the for the gTLDs, they are administered by a US-based company, it should be obvious to anyone registering in those TLDs that they are subject to US laws and regulations. Further, it isn't clear to me that the uniform dispute resolution policy is predominantly in favor of US based holders of US trademarks. This may be true for NSI administered TLDs, but is most likely not the case for non-US-based TLDs.

COOK Report: (to Conrad's response on the disaster represented by inconsistent responses from multiple roots) Under multiple roots some percentage of a businesses clients wouldn't be able to get where they wanted to go, but the majority of email would work if there were reasonable coordination would it not?

Conrad: Maybe. Maybe not. Depends on what happens to be in what caches and which caches a client DNS talks to and whether or not the primary is looking at the same roots as the secondaries and whether or not the connectivity to the roots is sufficient and if the root name servers can handle the load and [so on]

COOK Report: From your earlier remark about the "use of non-US-ASCII for domain names (e.g., daigaku.nihon (== ac.jp) encoded in Japanese), " i understand you to be saying that this WILL require a new gTLD for each non ascii alphabet?

Conrad: My personal preference would be for the creation of new ccTLDs, encoded in UTF-8 as CNAMEs off to the existing ccTLD. E.g., <utf-8 encoded nihon> CNAME jp. This would imply the current ccTLD administrators also manage the internationalized DNS version of their country's TLD, avoiding a whole raft of political nightmares.

With respect to gTLDs, I would be _very_ surprised if NSI is not looking at supporting internationalized DNS names within .COM/.NET/.ORG — the number of usable names in those domains is asymptotically reaching zero and opening up new character sets will result in a huge increase in available names (albeit to non-English speakers).

As to the creation of new i18n's gTLDs, I would imagine ICANN hasn't thought too much about it and if they do, I would guess they'd probably make a requirement that any new gTLD (and perhaps any existing gTLD too) should have i18n'd names CNAME'd off to the US-ASCII version (or, if they want to be really PC, allow the non-US-ASCII version to be the canonical name with the US-ASCII version CNAME'd to it).

The end result of all of this is what I call the babelization of the Internet. Unfortunate, in that it means we're about to embark on a deliberate disenfranchisement policy, however this merely reflects the real world: it is difficult for me to speak with a (say) Hakka only speaker... Oh yeah, and there is the minor problem that the use of UTF-8 will likely break most existing application protocol software (e.g., mailer user agents, netnews user agents, etc), Conrad concluded.

COOK Report: In addition to this online exchange we had an hour long conversation with a technical friend who is knowledgeable about DNS. This person pointed out while a split or splits in the root were certainly possible, unless there were many splits and they were of an extremely raucus nature, he thought the longevity of their impact would be short. To be of any global value, the root.cache file of the new root started by a single country would need to be picked up and carried by one or more of the 13 global root servers in order to be able to be seen outside national borders. He thought that achieving this agreement would be likely. But he also pointed out that because of the long time voluntary nature of the root server operations as a contribution to the "Internet community" that he believed that the root server administrators would cooperate among themselves to heal breaches in DNS connectivity caused by national defections from the DoC/ICANN root. He asserted that unless the episodes splitting the root were extremely violent, the impact on the net would be minimal because the breaches in DNS connectivity would be quickly ironed out. "The entire issue of stability goes to the good-will and agreement of the root operators. All it takes is one of them to announce the existence of the new country root's (possibly new) TLDs".

Some critics believe that this person's views reflect the world-view of the idealized, cooperative, pre commercial internet. They assert that DoC could quickly obtain a replacement for any operator of a global top-level root server that failed to cooperate. How effective this threat of "follow-our-orders-or-you-will-be-replaced" would be is untested. One reason that it is untested and a lot of people would like to leave it that way is that it is highly probable that many secondary and tertiary level servers would have to change their root.cache in order to be able to use a newly designated DoC/ICANN server

A side conversation with Bill Woodcock (which he gave us permission to quote), cleared up some issues regarding the permissible number of global root servers.

Woodcock: A quick note with regard to 13-root limit. The limit is on 13 _ip addresses_ of root servers in a UDP packet, not upon 13 _servers_. Multiple servers may share a single IP address and be topologically load-balanced. This also works to help overseas users, since they don't get stuck having to go transoceanic to hit roots all the time. I spoke on this at the Montreal IEPG (circa 1997), and I believe Bill Manning actually gave it a try. That work was based on some redundancy work we'd done previously for Oracle.

COOK Report: Could such servers be on different sides of a wan sharing the same IP address?

Woodcock: Yes, or more to the point, you could have an _instance_ of j.root-

servers.net in Tokyo, another _instance_ of j in London, and a third _instance_ of it in Cape Town. The routing tables would automatically direct you to the nearest instance, since they share the same IP address. Hypothetically. You'll need to talk with Bill Manning to determine how far deployment has actually gotten.

COOK Report: To sum up: it is hard to foresee with any certainty exactly what would happen if ICANN went away and the global e-commerce players were left to fend for their own interests vis-a-vis the root. Certainly the Internet would continue to work but with more uncertainty and less efficiency. Furthermore, the ongoing uncertainty of how much of its market and e-business would be reachable in the event of trouble would take the buzz off of most e-commerce business plans.

Conclusion

We are faced now with the very question that Einar Stefferud asked on May 13, 1984 nearly 16 years ago:

So, we can choose to spend our time sorting through this political morass without either the political savvy nor political clout to resolve anything, or we can go back to working on the technical syntactic structure within which the political decisions will eventually be made.

The main criterion for success will be whether we can invent a technical structure that will prevent political decisions from making a mess of Internet implementations. There just ain't no way that us techies are going to be allowed to dictate domain structures beyond the current bounds of the ARPA and DDN sub-nets.

Reset, Restart ... Stef

But since the root currently works, the easy way out has been ICANN, IETF and DOC to elevate the necessity of a monolithic root (rather than any diverse and locally accountable infrastructure) into a fatalistic doctrine designed to stifle any possible effort to engineer diversity. With all the eggs in the same basket a myriad of unforeseen perturbations could break them. To sum up the argument we have made, the level of fear that an alternate resolution model could emerge is very real. One classic scenario would be that the Chinese or some other authoritarian government could order their ISP's to point to a "approved" root structure is real. It also not too far fetched to see the US Government demand exclusion from the root for certain activities deemed totally illegal under US law like off shore betting. Plus ca change plus la meme choose.

In his thinking essay, Ed Gerck makes a call well worth heeding when he says: "It is my belief that the interests of small business, of big business, the public and that of governments would be better served by helping to seek a solution to the domain name issue, rather than by fighting the symptoms of the Domain Name Syndrome. The rhetoric around ICANN has been so intense though, and with such bad animus, that it requires a certain distance to see the forest for what it is.

[But how to get the distance?] "Readers may ask what is the schedule to propose new standards based on what I and my group are working on for domain names? As I see it, Gerck continues, time is not a trigger for the events needed to get us out of our predicament, but understanding is. Cooperation has its own dynamics and we must allow for things to gel, naturally. We can motivate, we can be proactive but we must not be dominating. We seek collaboration, not domination. Both technically as well as market-wise.

The issue is also user awareness. A perfect technical solution to a non-perceived problem is a perfect market flop. This [DNS as a single handle of control and point of failure] needs to be overcome with another design, under a more comprehensive principle, but one which must also be backward-compatible with the DNS."

Ed Gerck does not claim have the sole workable solution handed down to him from on high like the tablets to Moses on Mt. Sinai. To the contrary, he is defining a means of joining together to work our way out of this dilemma. It would be encouraging to see some interest in Washington in the incubation of the understanding necessary for the internet and e-commerce to cooperate in working its way out of the win-lose control situation in which it finds itself. The route of control has been tried. As we have shown in this discussion, not only has it not worked, it also looks to be untenable on a long term global basis.

May 2000 COOK Report

Commerce Department Formation of ICANN Seen as Illegal End Run Around the Administrative Procedures Act and the United States Constitution

Michael Froomkin's Findings to Be Published in Duke Law Journal - Lawrence Lessig Lauds Froomkin's Creation of Framework that Could Force Reform

Since the Internet is decentralized and controlled at the edge, the only tool offering the hope of affecting control of the entire Internet has been the DNS. Without a fully functional DNS the Internet is unusable. Control of the DNS hierarchy starts with control single global root relied upon by the 13 global root servers to enable users to get to the globally unique destinations they wish to reach. Control over the sole authoritative root is the only thing that in a broadly misunderstood effort to privatize the DNS, that the US government has not yet given to ICANN.

For almost four years the *COOK Report* has chronicled the unfolding struggle for control of the DNS and the root. A struggle that for the last two years has been focused on ICANN. The rise of ICANN is a story of hideous complexity. Most people without vested interests in ICANN's success have turned away in dismay as they observe its arbitrary acts, arrogant behavior and Byzantine structure. It leaves the lay person with a gut level sense that something is horribly wrong. But it is difficult for such a person to articulate with clarity precisely what is wrong and why it is wrong. The reason is that until now there has been no road map that identifies with precision what went wrong. At long last, University of Miami law professor Michael Froomkin has provided this road map in a landmark paper: "Wrong Turn in Cyperspace: Using ICANN to Route Around the APA and the Constitution" to be published by the *Duke University Law Journal*, October 2000, Volume 50, No. 1. The printer is expected to deliver the physical copies themselves before December 1, 2000. The article will be available of the Law Review web pages. The penultimate draft was placed at http://www.law.miami.edu/~froomkin/articles/icann1.pdf on October 16, 2000.

What follows is our attempt to distill the essential chain or arguments in this mammoth 166 page and 711 footnote long paper. We do so because we believe that the question of what happens to ICANN will be possibly the most important legal issue of the new decade. If this ICANN succeeds, there will be others to follow. Therefore we suggest that the determination of ICANN's fate is a decision that will impact whether we live in a nation state with a legal system accountable to individual citizens or a corporatist state that is run on behalf of the private profits of global corporations. If it is the latter we will find that we have given up the critical philosophical foundation on which our nation was built. That our government exists first and foremost to preserve and protect the rights of each of its citizens.

Instead of holding these rights sacred, our government has taken the example set by Esther Dyson and Mike Roberts in their "leadership" of ICANN to make clear the insidious nature of the attraction of the Clinton WTO mantra of self-regulation for those internet players involved in content and intellectual property. Here self-regulation is self-dealing and constant conflict of interest. Here the purpose of the self-regulation is to avoid government responsibility for the due process rights of the little guy. Self-regulation of the physical network has so far worked adequately. Howver, ICANN's self-regulation of the DNS on behalf of trademark holders is a self-serving attempt to pretend that conflict of interest can be wished away and that a century or more of American administrative law is no longer meaningful. The administration's self-regulation propaganda will continue to prevail only if the legal system of the United States begins to enfranchise only corporations and confiscates what used to be rights of citizens.

Froomkin has looked at the creation of ICANN and has carefully crafted an explanation of why the creation of ICANN by the Department of Commerce violates American law. This article will summarize his chain of reasoning. In so doing it will seek to disseminate more widely an articulation of why ICANN's ascendancy can and should be challenged. We believe that it is very important to use Froomkin's compelling insights to educate both citizens and the executive and legislative branches of the US government. We need to understand quickly what has happened and why we should "be afraid." Out of such education it is to be hoped that legal or legislative redress may be found.

ICANN is preparing to grant new gTLD franchises to powerful corporate and international agency applications. Once these are in place it is difficult to imagine a successful challenge to what has been done on ICANN's behalf. Individual and small business users of the internet need to be grateful to Professor Froomkin for his research and to the Duke University Law Review for its publication of this extraordinarily important article.

Editor's Note: This article is based on the draft of September 27, 2000. The article has been under going continuous revision since then. A check of most of our quoted passages shows no significant change between the draft of September 27 and that of October 16, which except for some cross correlation of footnotes to be done by the Duke Law Journal editors is final.]

The October issue of the Law Journal is devoted to the law school's thirtieth annual symposium on administrative law. This symposium was dedicated to understanding the issues presented by ICANN. In addition to the Froomkin article it contains an Introduction by Professor James Boyle of Duke University School of Law and a second article "ICANN and the Problem of Legitimacy by Professor Jonathan Weinberg of Wayne State University. When we asked Professor Weinberg for a very brief summary of his article he responded: "The article examines three ways in which ICANN has sought to establish its own legitimacy. It concludes that neither ICANN's adoption of the trappings of an administrative agency, nor its adoption of structures that make it resemble a representative (that is to say, elective) government body, nor its invocation of "Internet consensus," are successful in supplying legitimacy."

DoC ICANN Violates Either APA Or the Constitution

Froomkin's indictment in his opening paragraph is succinct: *"The United States government is managing a critical portion of the Internet's infrastructure in violation of the Administrative Procedures Act (APA) and the Constitution. For almost two years the*

Internet Corporation for Assigned Names and Numbers (ICANN) has been making domain name policy under contract with the Department of Commerce (DoC). ICANN is formally a private non-profit California corporation created, in response to a summoning by U.S. government officials, to take regulatory actions that the Department of Commerce was unable or unwilling to take directly. If the U.S. government is laundering its policy making through ICANN, it violates the APA; if ICANN is in fact independent, then the federal government's decision to have ICANN manage a public resource of such importance, and to allow - indeed, require - it to enforce regulatory conditions on users of that resource, violates the non-delegation doctrine of the U.S. Constitution. In either case, the relationship violates basic norms of due process and public policy designed to ensure that federal power is exercised responsibly."

He assumes that the Department of Commerce was serious when, in response to presidential directive in the summer of 1997, it talked about rule making. However when the Green Paper of late January 1998 was poorly received in Europe the Department of Commerce changed its course and authored what became known as the White Paper policy statement in early June 1998. Froomkin correctly characterizes this White Paper policy as: "Abandoning earlier hopes of issuing a substantive rule, which requires statutory authorization and is subject to judicial review, the policy statement instead set out goals that the administration thought could be achieved without rulemaking."

At the end of November of 1998, the Department of Commerce signed a Memorandum of Understanding with ICANN indicating the two would seek to work together to carry out the goals of the White Paper.

Evading Notice, Comment, Due Process and Judicial Review

One of ICANN's most significant acts has been its imposition of the Uniform Dispute Resolution Policy (UDRP). Froomkin points out that DoC has in some cases covered up what he later shows to be policy making by ICANN as if "they were matters of contract" falling "outside the rulemaking strictures of the Administrative Procedure Act." In other cases he finds that that the DoC treats ICANN as " an arm's-length private body exercising autonomous choices that could take effect spontaneously, without DoC's participation or responsibility. DoC has thus made, or acquiesced in ICANN's making, via research contracts rather than agency adjudication or rulemaking, some of the most important decisions relating to the near-term future of the Internet thus evading notice, comment, due process, and judicial review."

Froomkin suggests that arguments for privatization are often made on grounds of economic efficiency. In this case however "rather than privatizing a revenue-generating function, the government is privatizing a policy-generating function and because the privatization is subject to sufficient strings to make ICANN's actions fairly chargeable to the government." He correctly points out that not enough is known about the details of how much authority the DoC exerts to enable one to ascertain whether DoC tells ICANN what it must do in deciding matters of policy. If this is the case Froomkin finds that "DoC's use of ICANN to make rules violates the APA."

There follows a powerful paragraph that deserves quotation in full: "There is substantial evidence, discussed below, that DoC has directly instructed ICANN on policy matters. Furthermore, as ICANN is utterly dependent on DoC for ICANN's continuing authority, funding, and indeed its reason for being, it would be reasonable to conclude that the corporation is currently so captive that all of ICANN's decisions can fairly be charged to the government. If so, the DNS has not in fact been privatized at all, even temporarily. At least in cases where ICANN does what DoC tells it, and arguably in all cases, DoC's use of a private corporation to implement policy decisions represents an end run around the APA and the Constitution. To the extent that DoC launders its policy choices through a catspaw, the public's right to notice and meaningful comment, to accountable decision-making, to due process, and to protection against arbitrary and capricious policy choices, self-dealing or ex parte proceedings are all attenuated or eliminated; so too is the prospect of any meaningful judicial review. The result is precisely the type of illegitimate agency decision-making that modern administrative law claims to be most anxious to prevent."

Delegation of Public Power to Private Groups Ensures No Accountability

The only logical choice to make if one believes that DoC does not dictate policy decisions to ICANN is that it has decided to allow ICANN to make its own decisions. If this is the case Froomkin finds that DoC "violates an even more fundamental public policy against the arbitrary exercise of public power, the constitutional doctrine prohibiting the delegation of public power to private groups. Most famously expounded in two pre New Deal cases, Carter v. Carter Coal Co. and A.L.A. Schechter Poultry Corp. v. United States, the private non-delegation doctrine focuses on the dangers of arbitrariness, lack of due process, and self-dealing when private parties are given the use unaccountable private body that many feel has already abused its authority in at least small ways and is indisputably capable of abusing it in big ways, and nearly (but, as argued below, not quite) eliminated the possibilities for judicial review of critical decisions regarding the DNS. So far, ICANN appears to be accountable to no one except the DoC itself, a department with a strong vested interest in declaring its DNS 'privatization' policy to be a success."

ICANN suffers virtually no constraints in what it does. It is not constrained by the due process, public review and comment procedures designed to ensure the accountability of government administrative agencies. While the actions of corporations are restrained by competitors and shareholders, ICANN has no competitors and because it is non profit, no shareholders. While many non profits have members who can challenge misbehavior, ICANN has gone to great lengths to deny its members the redress that they would normally have available to them under the California public benefits corporation statute by which they are incorporated. While most non profits are restrained in their ability to raise funds, ICANN has turned its DNS registration franchises into sources of substantial income and with ICANN's recent fund-raising tactic of charging substantial fees for reviewing new top level domain applications, ICANN is getting even more funding.

Froomkin concludes: "The result is a body that to date has been subject to minimal accountability. Only the Department of Commerce (and in one special set of cases NSI or its registry) currently has the power to hold ICANN to account. NSI currently has no incentive to use its limited power and DoC has nothing to complain of so long as ICANN is executing the instructions set out in the White Paper. The accountability gap will get worse if DoC gives full control of the DNS to ICANN."

Froomkin examines in detail the question of whether ICANN is engaged in rule making in carrying out its DNS policy. He finds that " if ICANN is engaged in policy-making, and if DoC is reviewing these decisions, and retains the authority to countermand them, then DoC's adoption of or approval of ICANN's regulatory and policy decisions are subject the APA."

"If, on the other hand, ICANN is engaged in policy-making and DoC does not retain the power to countermand ICANN's decisions, then DoC has delegated rulemaking and policymaking power to ICANN. This probably violates the APA since it was done without a proper rulemaking; regardless of

the applicability of the APA, it violates the Due Process Clause and the non-delegation doctrine of the U.S. Constitution, as well as basic public policy norms designed to hold agencies and officials accountable for their use of public power."

Froomkin finds that "since ICANN's Board and staff operate largely in secret, it is difficult for outsiders to know how much influence the DoC has over ICANN's decision making. As a result, the statutory and constitutional arguments in this paper are presented in the alternative. The two arguments are very closely related, however, in that both rely on legal doctrines designed to promote accountability and prevent the arbitrary exercise of government power."

The DoC ICANN Defence: Private Party

Froomkin explains how DoC and ICANN defend themselves against his charges: DoC claims that ICANN is a private party operating at arm's length removal from DoC and that ICANN is a standards setting rather than a policy making body. Government agencies are forced by the Administrative Procedure Act to pay attention to due process concerns designed to see that citizens do not suffer from the arbitrary use of government power. DoC argues that because ICANN is a private corporation it is not bound by the APA.

Froomkin's rebuttal is vigorous and correct: "In fact, as detailed below, ICANN's relationship to DoC is nothing like the arm's length relationship suggested by the private party story. Although ICANN is private, it is no ordinary corporation, and its relationship with DoC is highly unusual. ICANN is totally beholden to DoC for its creation, its initial policies, and especially DoC's loan of control over the root. This control over the root is the sole basis of ICANN's relevance, power and financing, and DoC can take it away on 120 days notice. More than anything, ICANN seeks to achieve permanent and perhaps irrevocable control of the root when the current MoU expires. DoC has some control over ICANN through the stick of the MoU, but the real control comes from the carrot. ICANNs ability to retain or expand its control over the root is entirely at DoC's discretion."

Froomkin explains three reasons why DoC cannot be considered to have an arm's length relation ship with ICANN. First ICANN does not have the authority to create a new TLD without DoC approval. Neither NSI may change the content of the root without written DoC approval. "Second, DoC's retention of the power to take control of the root away from ICANN is of enormous significance, as it forces ICANN to be exquisitely conscious of DoC's requirements. If ICANN fails to meet DoC's expectations, DoC can choose another body to replace ICANN. All of the relevant agreements provide that if DoC recognizes another entity as 'NewCo" in ICANN's place, then the obligations to ICANN in those agreements immediately terminate. "Froomkin continues "ICANN's only reason for existence, and the sole source of its power over the DNS is that the thirteen root servers treat it as authoritative, and that the government instructed NSI, another contractor, to both defer to ICANN's policy and to pay it money . The root servers recognize ICANN only because DoC signed an MoU with ICANN and announced that ICANN is the relevant authority. Were the U.S. government to transfer its recognition to another authority, the root servers would be under no more legal obligation to recognize that new authority than they were to recognize ICANN, but their move is all but certain. So long as the root servers recognize ICANN's authority, ICANN is able to dictate contractual terms to registries and registrars who wish to be included in the legacy root. By dictating terms to registrars, ICANN can also enforce terms on registrants, since ICANN can (and does) require that registrars include standard terms in contracts with registrants."

"Third, DoC and ICANN have a warm and cooperative relationship, although whether that relationship is best characterized as a partnership, master-servant, or self-regulatory body and supervising agency is hard to discern. Whatever the precise nature of the relationship, it certainly is not arm's-length. In particular, it cannot reasonably be characterized as DoC calling ICANN into being, signing an MoU with it to give it authority, and letting it go off on its own."

Standards Making

Since ICANN is not a private arm's length entity free to adopt its own policy, Froomkin next explains why Dyson, Roberts and Sims always go to such lengths to claim that ICANN does nothing more than technical coordination. If this were true then ICANN would be only a standards setting body and Froomkin points out that there is substantial legal precedent for the federal government to grant standards setting authority to private corporations. As Froomkin puts it those who defend ICANN say: "ICANN does not do 'policy'; if there was any policy to be done (DoC is a little vague on this), it was done in the White Paper-a statement of policy. And ICANN most certainly does not do 'regulation' or 'governance'. ICANN is at most implementing the key pieces of the White Paper policy - privatization, Internet stability, increasing competition, bottom-up coordination." He continues: "if it were true that ICANN was limited to "technical coordination" that would rebut the claim of an unconstitutional delegation of power. In fact, as detailed below, the standard setting story ignores reality. While some of what ICANN does can fairly be characterized as standard-setting, key decisions would certainly have been rule making if done directly by DoC, and remain regulatory even when conducted by its proxy [ICANN]."

Avoiding All Accountability

Later in his paper Froomkin finds that "DoC is not the first agency to seek to use the corporate form or to create a private corporation to achieve desired ends. The Government Corporation Control Act (GCCA) is Congresses most comprehensive modern attempt to define when and how federal officials may use private corporations for public purposes." He then notes that corporations by the executive branch 'to act as an agency' without specific legal authorization." A page later Froomkin notes that the GAO's General Counsel in its examination of the FCC's creation of the Schools and Libraries Corporation and the Rural Heath Care Corporation dismissed the idea that either organization was independent of the FCC by stating that the GCCA "prohibits an agency from creating or causing the creator of a corporation to carry out government programs without explicit statutory authorization." In short one way that ICANN could be created outside the constraints placed by the Government Corporation Control Act was for the DoC to decide to publish the White Paper as a "formally non-binding policy statement." Doing this appeared "to be a successful end-run around the GCCA - which suggests that the GCCA may need amendment." For Froomkin concludes: "By calling for NewCo to form 'spontaneously' government officials avoided directly creating the corporation." In doing so they avoided the accountability that the GCCA would have brought to ICANN.

Under a heading called "Structural Failures / self-dealing" and a sub heading of"Time" the Froomkin draft states: "ICANN's structure taxes time as well as money. The profusion of constituencies, working groups, ad hoc committees and the like means that only those with an enormous amount of time to devote to ICANN issues can stay abreast of every developing 'consensus' policy. In practice, those who can afford to pay someone to represent them - predominantly commercial interests who hire lawyers or delegate managers to be their spokespersons - are able to dominate. Un-organized groups such as users or small businesses must rely on volunteers, and tend to be outnumbered in committee."

On the following page Froomkin adds: "For an entity designed to make rules relating to the Internet, ICANN and its subsidiaries seem oddly dependant on attendance at physical meetings. Because ICANN sees itself as global, its meetings are peripatetic. This policy, which has the advantage of making it more possible for geographically disparate groups to attend the occasional Board meeting also makes consistent participation by those without substantial expense accounts impossible. People who cannot attend meetings of the Board are able to take part in debates in only the most limited and derivative way - the few remote comments read to the meeting are filtered and edited by the readers, as contrasted to attendees who queue for a microphone and can say whatever they like." So much for the public notice and comment rule-making procedures of the APA.

Beware— More ICANNs to Come

"The issue then is whether an attempt to vest power in an agency or a private body is constitutional. In the case of ICANN, there is no statute. Congress at no time determined that the DNS should be privatized, or indeed legislated anything about national DNS policy. Instead, DoC itself chose to delegate the DNS functions to ICANN, relying on its general authority to enter into contracts. ICANN is also a very unusual corporation. There are many government contractors, both profit-making and non-profit. But it is unusual for a nonprofit corporation to be created for the express purpose of taking over a government policy-making function."

"There is a danger, however, that ICANN may not be unique for long. One administration spokesman has already suggested that ICANN should be a model for regulation of other Internet-related issues such as accreditation standards for distance learning, and e-commerce over busness-to-business "closed"networks." [Froomkin cites: DoC General Counsel Andrew Pincus, Remarks at the Meeting of the Bar of the District of Columbia (Apr. 25, 2000); see also Brian Krebs, Regulators Would Do Well To Mimic ICANN - White House (visited Aug. 18, 2000) <htip://www.newsbytes.com/pubNews/00/148011.htm1>] Froomkin continues; "The specter of a series of ICANN clones in the U.S. or in cyberspace should give one pause, because ICANN is a very bad model, one that undermines the process values that motivate both the APA and the due process clause of the Constitution."

Toward the end of his paper Froomkin returns to this theme: "For all of its problems, ICANN commands substantial support, and not all of it from the stakeholders who dominate it. Many believe that the need for some sort of management of the DNS is so pressing that, until they see a viable alternative, ICANN is the only game in town (others, who wish to preserve the status quo as long as possible, may see ICANN as the least bad way of blocking change). To many of the participants in the DNS wars, especially old ICANN hands imbued with the libertarian traditions of the founders of the Intemet, anything that takes power away from government is presumptively, and sometimes irrebuttably, better than a governmental solution."

"These views are mistaken, if only because there is more at stake here than the Internet. Even if ICANN were thought to be a good thing, [seeing it just in the context of] a narrow focus on the Internet ignores the pernicious effect of ICANN on the U.S. government itself and on our democracy for there is a real danger that ICANN will not be a fluke, but will be used a model for additional erosions of responsible government. DoC's use of ICANN undermines accountability. Every government power must be exercised in accordance with law and with the Constitution. But ICANN is a private non-profit California corporation; unless it is a government actor or advisory committee, neither the APA nor the Constitution apply to it.

The APA and the Constitution apply to DoC, however, and this is where the main violation of law is to be found. Allowing DoC to use ICANN to make non-technical policy decisions violates basic norms of governance and accountability. DoC cannot quasi-privatize the DNS in a manner that allows the US to retain ultimate control of the root zone file but achieve deniability about everything that its agent or delegate does with day to day control. Depending on one's view of the DoC-ICANN relationship, the violation may be a failure to make a general and prospective rule in conformity with the APA's procedural requirements, or it may be the delegation of so much control to ICANN that it violates the Carter Coal doctrine [of non-delegation of public functions without congressional authorization] under the guise of technical standard setting." Additionally, the DoC-ICANN alliance is an endrun by the federal executive branch which circumvents the role and oversight of the federal congressional and federal judiciary functions.

ICANN's Choice of New TLD's Not a Technical Matter

Given that ICANN has received 47 applications for more than 100 new TLDs ICANN is likely to find that as Froomkin says "there are more technically qualified applicants for new gTLDS than ICANN wishes to create, and ICANN chooses among them based on some idea of the 'quality' or 'usefulness' of the proposed gTLD, or the extent to which the applicant promises to enact social policies such as privacy or trademark protection, then ICANN will have clearly crossed the line into making social rather than technical policy. It cannot seriously be suggested that choosing whether the world is better served by".banc," ".shop," or ".xxx" is an issue with any "technical" content whatsoever. This is a question of social and political import only, the technical issues are identical whatever a TLD happens to be called. As such, the selection among technically qualified applicants on social merit is an adjudication that DoC cannot delegate to ICANN so long as the DNS remains in DoC's ultimate control."

Froomkin shows some approaches that might be made toward correcting the travesty of justice that ICANN has become. He concludes: "If the current quasi-privatization of the DNS is illegal, then DoC needs a new policy to deal with the DNS. (Indeed, I am certain that uncertainty as to how to answer the question, "If not ICANN, then what?" is an important part of the reason why opposition to ICANN has yet to reach critical mass.) DoC has at least four options."

Having promised Professor Froomkin that we'd stop here we refer our readers to the concluding pages of his article - available we hope in about 30 days.

Froomkin's final call is appropriately eloquent. We can only wonder who will hear it? "The ultimate problem with DoC's reliance on ICANN is not the corporation's secretiveness, or its decisions, or its fight against accountability. Rather the problem is that DoC's reliance on ICANN, and endorsement of its activities, undermines basic elements of accountability and due process on which the administrative state is based. By lending ICANN its control over the DNS, DoC created a system in which social policy is made not by due process of law, but by something that begins to resemble government-sponsored extortion."

Lawrence Lessig Comments

In a brief interview on October 17, 2000 Stanford Law School Professor and ex candidate for the ICANN board Larry Lessig complimented Professor Foomkin for having drawn a well argued road map of grounds for legal action against the US Department of Commerce for its actions in creating ICANN. Lessig said that with the legal research solidly done it would now be much easier to convince those opposed to ICANN to unite around a legal challenge that would most likely be class action in form.

Lessig said that he would prefer an ICANN

done right, namely one that did only technical coordination and not policy, to a Department of Commerce entity that did make policy subject to what he regards as a cumbersome, bureaucratic and expensive Administrative P rocedure Act. The question of course becomes how to convince ICANN and DoC to change their agenda.. To this end a group ready to bring suit could approach with demands. If the demands were not met, the group could file suit. One way of obtaining a change in DoC ICANN behavior could be a settlement that is put in place as a condition for dropping the suit. Faced with an unambiguous certainty of up to three years of litigation, the forces behind DoC's and ICANN's position would listen to reason rather than have the suit go forward, or so Lessig hopes.

ICANN Member at Large Elections

On October 11, it was heartening to see Kal Auerbach and Andy Müller-Maguhn win member at Large board seats. Also on October 11 a BWG member wote: "There were 3 different elections:

1) The marginals: Africa & Latin America: Players ride. 130 votes cast in Africa; 1,402 in L.A. In both ICANN nominated candidates got over 80% of the votes.

Note however that the person elected from Africa, although not an ICANN critic at all, is very fair minded and supports the GA's role and the role of individuals. I don't know much about Campos, but his replies on the forms suggest he thinks ICANN is basically ok, not too many new TLDs needed, but UDRP and other reforms are.

2) The rich West: Not happy with ICANN.

N.A. Almost a marginal. Only 3,449 votes cast. Only 32% for ICANN nominated candidates, including Lessig. Almost 3/4 voted to repudiate current icann policy if you just count count Simons, Lessig & Auerbach. The anti-icann crowd beat the pro-icann crowd. Auerbach and Lessig were ALLIES and cooperated in slating with Simons; together they got 75% of the first choice ballots in N.America.

Europe. A third of the total ballots. Only about 1/4 for ICANN nominated candiates. Both victors are ICANN critics. Expect fireworks.

3) Asia. The big one: A vote for ICANN status quo. 17,745 votes - more than half the total cast. 95% for ICANN nominess. The winner says "In my opinion, ICANN has interpreted its mission in a cautious but still proper manner. ICANN has a coordination function to perform to ensure the seamless operations of the Internet. In exercising its coordination function, ICANN must take into account the legal and policy impacts of its management decisions, and therefore its scope of mission is not limited to purely technical issues." Thinks the UDRP is at least OK, maybe great.

Winner is resident in DC area! Number one goal listed at http://www.mkatoh.net/FAQ.htm#Goals is: "To make ICANN a 21st Century model for international organizations" See also http://www.mkatoh.net/speech/icann_katoh072000-e.ppt (arguing ICANN should be model for 21st century organizations). More than half the voters are in the Asia/Pacific region. And they picked a DC IP lawyer [who is a Japanese citizen and an employee of Fujitsu]."

The October issue of the Law Journal is devoted to the law school's thirtieth annual symposium on administrative law. This symposium was dedicated to understanding the issues presented by ICANN. In addition to the Froomkin article it contains an Introduction by Professor James Boyle of Duke University School of Law and a second article "ICANN and the Problem of Legitimacy by Professor Jonathan Weinberg of Wayne State University. When we asked Professor Weinberg for a very brief summary of his article he responded: "The article examines three ways in which ICANN has sought to establish its own legitimacy. It concludes that neither ICANN's adoption of the trappings of an administrative agency, nor its adoption of structures that make it resemble a representative (that is to say, elective) government body, nor its invocation of "Internet consensus," are successful in supplying legitimacy."

December 2000 COOK Report

Light, IP and Gigabit Ethernet: The COOK Report on Internet Annual Report 2001

Introduction

Part One

The Internet Way **p. 3**

Part Two

National Trials of "True Internet" **p. 50**

Part Three

Toward a Bandwidth Commodity Market **p. 126**

Part Five

Other Oncoming Technical Developments that Devour Bandwidth