FEDERAL TRADE COMMISSION

I N D E X

FEDERAL TRADE COMMISSION

In the Matter of:)
EMAIL AUTHENTICATION SUMMIT)
a corporation.) Matter No. P044411
----------)

TUESDAY
NOVEMBER 9, 2004

Federal Trade Commission
601 New Jersey Avenue, N.W.
Washington, D.C. 20001

The above-entitled matter began pursuant to notice, at 8:30 a.m.

P R O C E E D I N G S

MS. COLEMAN: Hello, and good morning to everyone. Yes, thank you all for being here so bright and early. We really appreciate this turn out. It's one thing to see a vision and to have an idea and then to see itself manifest itself into all of the faces today, so it's really a pleasure and an honor for us to have you all here at this very important Email Authentication Summit.

We want to go ahead and get started, and I'm going to introduce to you the chairman of the Federal Trade Commission, Deborah Platt Majoras, who will start the Summit off today by giving us a warm welcome and opening remarks.

Chairman Majoras was sworn in on August 16, 2004, as the chairman of the Federal Trade Commission. She joined the FTC from the law firm of Jones, Day in Washington, D.C., where she served as a partner in the antitrust law division.

I am pleased to introduce to you now Chairman Deborah Platt Majoras.

(Applause.)

CHAIRMAN MAJORAS: Well, good morning. I never expected this turn-out at 8:30. Maybe by 9:30 or so, so I'm really thrilled to see you all here so bright and

early, and on behalf of the Commission and our co-sponsor, the Department of Commerce, National Institute of Standards and Technology, I welcome you to this two-day Email Authentication Summit.

Currently, there's probably no more intractable consumer issue than spam. Spam poses two principal threats to electronic communications over the Internet for consumers and businesses alike. First, deception and fraud characterize a significant amount of spam. Indeed, spam apparently is the vehicle of choice for many deceptive and fraudulent marketers.

Second, spam, even if not deceptive, may lead to disruptions, inefficiencies and security breaches in Internet services. Spam often spreads viruses that wreak havoc for consumer users. Moreover, the sheer volume of spam now being sent is creating Internet infrastructure problems.

These problems impose significant costs on consumers and businesses and, importantly, threaten their confidence in the Internet as a medium for commerce and communication.

The FTC has pursued a threefold strategy to combat spam: Enforcement, education and research. We've brought nearly 65 spam related cases against some 165 individuals and firms, and we have worked very hard

to educate consumers and businesses about the risks from spam and how those risks can be combated, but as you know, your government cannot alone solve this problem.

Last spring the Commission held a highly successful three-day Public Forum that examined spam from all viewpoints. The Commission convened the Forum to learn more about the issues spam poses and to act as a potential catalyst for solutions to spam problems, brought together representatives from as many sides of the issue as possible to explore and encourage progress for possible solutions to the detrimental effects of spam. Today, in partnership with NIST, we continue those efforts by convening this Summit.

The Commission first raised the issue of authentication last June in our report to Congress on the possible creation of a Do Not Email Registry. The Commission concluded that without a system in place to authenticate the origin of email messages, a Do Not Email Registry not only would fail to reduce the burdens of spam, but in fact could actually increase the volume of spam sent, as illegal marketers might use the registry as a directory of legitimate email addresses.

Instead, the report recognized that solving the spam problem must begin with the recognition that spammers are essentially anonymous. The current email

system enables spammers to hide their tracks, thereby evading ISP's anti-spam filters and evading law enforcement. This is not a problem that lends itself well to governmental solution. The best hope is for the marketplace to develop and employ technological solutions to prevent spammers from hiding behind a technological veil.

In response, ISPs and others involved in the email system have proposed domain level authentication systems, systems that would enable a receiving mail server to verify that an email message actually came from the sender's domain; in other words, if a message claimed to be from ABC@ftc.gov, these private market authentication proposals, which you'll hear more about today, would authenticate that the message came from the domain ftc.gov. Now, it would not, however, authenticate that the message came from the particular email address, that is ABC.

Domain level authentication by itself will not solve the spam problem. It can, however, significantly impede spammers who engage in spoofing, the falsification of email headers, and criminals known as phishers, those who send emails that look like official correspondence from a financial institution and deceptively lure consumers into providing account

information that they then use to steal from the account holder.

Domain level authentication can also help ISPs and other operators of receiving mail servers reduce the incidents of false positives, that is legitimate messages wrongly identified as spam by spam filters. Domain level authentication can also enable the government and ISPs to identify more effectively, and then in our case, prosecute spammers who violate the Can Spam Act or other statutes.

The Commission's Do Not Email Registry report laid out a multistep process aimed at promoting wide scale adoption of domain level authentication systems. The first step in that process is today's Summit, in which the Commission and NIST have convened an impressive array of technologists to explore the nuts and bolts of various proposed authentication systems and to determine the necessary steps to achieve rapid deployment of email authentication, and I thank all of our distinguished panelists for your participation.

During today's sessions, we will receive a technological overview about email authentication and how it works. We'll also learn more about the technological basis for many of the industry email authentication proposals and the status in testing and

implementing these proposals.

Tomorrow, we will explore weaknesses that may exist in any of the proposals and how industry participants can expect to overcome these weaknesses. We will learn about what real world impact authentication will have and how this impact could ripple throughout the global community.

We'll learn how participants in the email arena plan to implement systems, and finally we'll hear about other services, such as reputation and accreditation services that may be required to render an email authentication system more effective.

We at the Commission, together with NIST, are pleased to provide a forum for discussion of the intricacies of domain level authentication. It is an important step forward, but talking about authentication will not be enough. As Ralph Waldo Emerson said: "Good thoughts are no better than good dreams unless they be executed."

After tomorrow, we urge you to take the knowledge you have and the knowledge you will have gained at this Summit and continue the process of making email authentication a reality.

The risk that spam will stymie realization of the Internet's benefits to consumers and the global

economy is too great to ignore and there is no time to waste.

Again, I welcome you, and I thank you, and now I'll turn the workshop over to the first panel. Thank you very much.

(Applause.)

"BACK TO BASICS: WHAT IS EMAIL AUTHENTICATION AND HOW DOES IT WORK?"

PARTICIPANTS:

SHERYL DREXLER, Investigator, Division of Marketing Practices, FTC

JOHN R. LEVINE, Taughannock Networks

MS. DREXLER: Good morning, everyone. I'm Sheryl Drexler. Thank you very much, Chairman, and we wanted to start with just a few brief housekeeping announcements, so bear with me a minute here.

First, if you have a cell phone or any other device that beeps, please, please, please turn it off. We also want to say in the event of an emergency, should there be one, which we don't expect there to be, but just in case, you'll be instructed where to go. Remember the exits are behind you and out to the front where you came in.

We wanted to thank Verisign for providing refreshments for the break this morning, and we also wanted to thank in advance the Direct Marketing Association, the Association of Interactive Marketing for providing refreshments on Wednesday morning and Cisco Systems Inc., is providing refreshments for tomorrow afternoon. There are trash cans out in the hallway for your convenience, so please use them.

We want to make sure that everyone on the panel speaks into the microphones so that people can hear, both on the phone as well as in the room, and, panelists, if you have something to say, you can turn your table tents upright and turn it back down to the horizontal position when you're done speaking.

We do want a lot of audience participation, and so when we do have questions and answers from the audience, we do ask that you wait for a roving microphone to reach you. Otherwise again people on the other side of the room as well as on the phone will be unable to hear you, and if you could also spell your name, your last name, and introduce yourself when you are asking the question.

For those people who are on the phone listening, if you would like to email questions to us, you can do so at Email Summit underscore Nov, as in November, 04

@ftc.gov. If you are a panelist or an audience member, you should hang on to your name tag throughout the day. Panelists, you want to hold on to yours throughout the duration of the Summit.

If you go out to lunch, bring your name tags with you. Otherwise when you come back in you'll have to get new ones. Whether or not you're a panelist or an audience member, you will have to go through security again, so please leave enough time to get through security when you come back from lunch. Remember seating is on a first come, first serve basis.

Now that we have all those announcements out of the way, we wanted to get started with the first panel. John Levine has been writing and consulting on email and the Internet for over a decade, and he's the primary author for the best selling "Internet for Dummies" and many other books. He's a board member of the Coalition Against Unsolicited Email, and since 2003 he's chaired the Anti-Spam Research Group.

It's now my pleasure to introduce to you John Levine.

(Applause.)

MR. LEVINE: Thank you very much, and thank you for inviting me to be the first panelist, and now I have to see if I can find my slides.

I apologize. Bear with me a second. I have a few different versions of my slides. I want to see if I have the right one.

With that I'm not going to attempt to give you a 40,000 foot concord's eye view of the email authentication issue, and I think some sort of the principles you should use to think about the various proposals and the various issues that are brought up during the upcoming two days.

I want to start by backing way up and saying, Why is email important? Why are we all here? Why do we care about email? The important thing about email is that it goes all over the world, and over the past 20 years using IETF standard email, we've managed to take what used to be little local email systems, put them all together into one global system.

So now we absolutely take it for granted that I, on my funky little network in upstate New York, can send email to any of you, and it doesn't matter whether you're on a big commercial provider like AOL or you're on a government agency here like the FTC or you're in a corporation like IBM, or you're somewhere in Europe or in Asia. It all just works, and although we take it for granted, it took a lot of work to get to the point where everything just works.

As we continue to evolve the email system, it's important to continue that and that it continues just to work because part of the process of authentication is a reversal of basically everything we've done over the past 20 years.

What we've done so far is to make it possible to send email from absolutely anybody to absolutely anyone else, and one of the things that authentication does is we're going to say there are some kinds of emails we don't want, so that the general theory of any sort of email authentication scheme is that we figure out which mail is good, somehow, whether signatures or source identification or any of the other dozen plans and acronyms that you're going to be hearing about over the next couple of days.

Okay. Here's all the mail, and if you can see the slides, the stuff that's in green, this is all the mail that we figured out must be good mail, so then here in red, this is all the mail we've all figured out must be bad mail, and depending on the scheme, either we've specifically figure that it's bad or we took out all the good stuff and what's left over must be bad. You say, ah-ha, now that there we know what the bad mail is, zap, we're going to get rid of it.

So once we have gotten rid of all the bad mail,

then presumably all that's left is all the good stuff, and the spammers will all go away, and we'll have our land of peace and plenty, right?

Well, sort of. The problem is that no matter what scheme we do, there's always some risk it's going to make a mistake, and so here I think this is the realistic prospect, which is most of the mail is identified correctly, but some of the mail isn't. Here some of the bad mail has been identified as good and some of the good mail is identified as bad, and no matter how wonderful the scheme is, there's always going to be some of that.

What we need to figure out is both how much of that is going to happen and how much can we put up with.

Now, there are I think four approaches to mail authentication, and you can tell this is a new field because they all have long, hard to pronounce, practically interchangeable names, but I'm going to attempt to divide the four general approaches into authentication, authorization, accreditation and reputation, and I'm sure there are people who will up and down and say I've defined them wrong, but bear with me because I think these are still four useful categorizations.

Authentication is this mail really did come from

so and so, or this mail really did come from so and so's domain, and there's a variety of schemes to do this, and again I'm not going to get into which ones do it, but authentication says, okay, this mail really is from Fred.

Authorization is back office stage. It doesn't say who this mail is particularly from, but it says, okay, if the mail came from this computer, then it could be from Fred, or it may just be that, well, if this mail came from this computer, then it's probably valid since there's some schemes that simply observe that some of the computers on the Net send valid mail, and most of the computers on the Net don't, so this case tries to sort of separate the sources, is this source authorized to send mail that is valid or some definition of valid.

Now, once we have started to separate them like that, it is way too hard for every possible recipient to make its only list of good guys and bad guys, so we're doubtless going to see accreditation schemes, which are basically senders come in and say or senders come in and prove their virtue, and basically an accreditor will say, These are people you can trust to send you legitimate email, but it's at the initiative of senders.

The flipside of accreditation is reputation.

All right. We got this mail from foo.com, never heard of them, are they any good? So you can go and we're all positing that there will exist things called reputation systems, although in fact none of them really exist yes, and the idea is you can go to the reputation system and say, hey, I got this mail from so and so, and it will come back with some sort of answer, like it might just say it's good, it's bad or it might say well, we've had 16 reports of good messages and 3,000 reports of bad messages or something like that, but reputation schemes are entirely up in the air.

Wearing my Anti-Spam Research Group hat, I've been attempting to crank up some research and reputation systems with surprisingly little success so far.

So we're going to do these four things, and if we're not careful, we're going to get into trouble because I see three related issues. First is the email world is really big and surprisingly fragile. There's all sorts of things that you could do that seem to be tiny to you, but in fact the mail would come grinding to a halt, and in particular, taking a system that's not designed to be secure and making it secure is really hard.

And a good analogy in this case is actually the postal mail system. There's lots of ways that the

postal mail system is not like the email system, but one way that they're absolutely the same is that they're both really large and they both process vast amounts of traffic, and neither one has a security model.

If I were mad at you, I could right your name on an envelope, and I could drop it into a mailbox, and that would be that, and the Post Office neither knows nor cares that it wasn't you that sent that message. We have unfortunately in recent years been forced to try to make the postal mail system somewhat more resistant to fraud and bad guys.

The example here in Washington, when the nutty guy in New Jersey was sending letters around with Anthrax. You think of the approaches they tried to do simply to make the mail around here more resistant to email and people mailing letters full of poison? What happened to the mail? It ground to a halt.

Partly it was because they had to apply them in a hurry, and they didn't have time to design something really good, but I would argue there's no way you can design something really good. Making an insecure system secure is really hard, so anybody who says, well, let's do so and so I would argue has not thought about the problem.

The third issue is what it says here, one man's

security hole is another man's handy facility, and there are some things that are unusual but legitimate. For example, when I'm sending email, nearly all of the mail I sent, I send through my mail server at home since that's the normal place I send mail. I don't always. I might be here, and I might be sending mail through a mail server at the Hilton if that's where the conference is.

The same thing with paper mail. If we wanted to make it -- imagine we were doing the same to paper mail, we wanted to make it so that any mail sent with my return address on it was actually from me. Well, normally I send mail from my own Post Office, and normally I mail it myself but sometimes I don't. Sometimes my wife mails it or sometimes I'm visiting my sister, and I might either mail the mail at her Post Office or she might send mail on my behalf at her Post Office.

You can come up with this long list of less usual, perfectly legitimate ways that I might send mail, and the exact same analogy applies in the email world. If you come up with all the ways you think people might legitimately send emails, and you will find no matter how hard you look, your list is not complete. There are legitimate ways of sending email that none of us have

thought of, and as soon as we make some sort of security system or authorization system that assumes everybody will do one of these six things, then we'll find the other 40 things people are doing, and we've broken their mail.

So what do we do? The Internet started as a research experiment, and to some extent it still is a research experiment, so we have to do lots of experiments. A message I hope we'll take away today is we have all sorts of really interesting proposals for mail authentication and mail security, and none of them are ready for prime time yet because before we can use any of them, we need serious, large scale experiments to find out how well they work, how expensive they are, how hard they are to maintain and what breaks, and we find stuff that breaks, then we have to come back and do it sort of jointly, as an Internet community, make a decision. Are we willing to put up with having something that used to work not work or do we have to go back and say we're going to try a different security approach that allows this particular thing to continue.

I can easily see situations where you might decide either but you can't just waive it off. It will be an issue.

The second thing is we have to have experiments

that go along multiple providers. I've done all sorts of little experiments on my tiny network at home, which I find fascinating, but I suspect would not be pervasive to say the AOL Postmaster, much so he may respect me, and any useful approach can only be useful if -- it has to be workable for everybody, all the big networks in the U.S., all the little networks in the U.S. and all the big and little networks in Asia and in Europe and in Africa.

If we have an authentication system that can't be used by somebody in a rural village in Africa at the bottom of a satellite link, we failed, because the Internet to people like that is one of the most important things the Internet does, and if we cut them off, we've done a vast disservice to them and to us.

This means as a result the proprietary approach simply can't work. Any approach that says, well, you have to use our proprietary stuff isn't going to work because everybody is not going to use it. It won't work unless it can work for everybody.

Finally, are we looking at a single approach? No, we were not. If we had a magic bullet, we would have shot it already, but we don't. Pretty much every approach I've seen proposed, certainly all the ones that people are going to describe today, can coexist. We can

do experiments with all of them at the same time. I'm simultaneously experimenting with signing my name and looking at the source authentication and doing various cryptographic things to check the return address.

I can do them all at once, and certainly for experiments we can do them all at once, and in practice we're probably going to do several of them at once because first we need to try them all in parallel and keep the ones that look promising, but more importantly, the bad guys are going to counterattack.

If we put all of our eggs in one basket, it means those guys are going to stomp on that basket. If you have multiple security approaches, then the chances of the bad guy circumventing all of the security approaches at once is much less. This is a familiar message from physical security, and it applies exactly the same way to computer security.

Many of us are here wearing badges with three or four letter acronyms on them, and I'm going to suggest roles that we all need to look to be playing in our various organizational roles. Software developers need to be developing the possible approaches and rolling them out, and in fact we've been doing a pretty good job at that. There are tests now of Sender ID, SPF and DomainKeys and Internet Identified Mail and probably

more if I thought about it.

The ISPs and network operators are starting to be very cooperative in trying them out, and what I have not yet heard back is reports on how well they work, but I think they will start to come back, and it is important to share results, so we can compare and say, well, if it works really well for one ISP and not for another, what are they doing differently.

The various standards organizations, the IETF and ITU, standards organizations are not good at developing technology. They're really good at codifying technology. I mean, once we have something that seems to be working, standards organizations are enormously helpful to actually nail down the details so that if I implement it or you implement it, it will work, and you'll say, well, gee, don't you expect this to work, ha. In writing a spec that actually clearly gets all the details correct is enormously difficult.

These are the areas where the IETF and ITU have considerable expertise, and the ITU also I think can provide political cover. They can go and advise their various member countries that this is not a plot by corporations that are going to kick them off the Net, and this really is appropriate technology for countries all over the world.

The FTC here can keep us honest and remind us there are laws that we have to comply with, and more importantly can document where law and technology meet. There are anti-fraud laws. Particularly there are laws about fraud related to spam. I was the expert witness in the Leesburg case two weeks ago that appears for the first time will put a spammer in jail.

Partly what we had to do was we had to say, this guy was doing these things which broke that law. Being able to codify that these authentication schemes are a common use, and if you break them, that's prima facie evidence that you're breaking the law. That's very useful, for making the laws more enforceable.

So here's my prescription for the next few days. The developers need to build a software. The network operators and the bulk mailers and the bulk recipients need to do experiments, and we all need to report and compare results. Standards organizations then need to help us get together and codify and standardize the results and get going and use it, so let's get going.

Thank you.

(Applause.)

PANEL 1: DEFINING THE FRAMEWORK: POLICY CONSIDERATIONS FOR EMAIL AUTHENTICATION

MODERATOR: COLLEEN B. ROBBINS, STAFF ATTORNEY, FTC

PANELISTS:

DUANE L. BERLIN, Lev & Berlin

SCOTT BRANDER, Harvard University

PAULA BRUENING, Center for Democracy and Technology

RAY EVERETT-CHURCH, ePrivacy Consulting

FRANK GORMAN, Bryan Cave, LLP

DAVID KAEFER, Microsoft Corporation

ANNALEE NEWITZ, Electronic Frontier Foundation

DANIEL QUINLAN, Apache SpamAssassin, Apache Software Foundation

JONATHAN ZUCK, The Association for Competitive Technology

MS. ROBBINS: Good morning. All the panelists for Defining the Framework please take your seat up at the front table.

Good morning. My name is Colleen Robbins, and I'm an attorney here with the Federal Trade Commission in Washington, D.C. Welcome to this morning's panel on Defining the Framework: Policy Considerations for Email Authentication.

This will be a discussion about various policy

and legal issues as they relate to email authentication, and the individuals who are going to address these issues are as follows: Starting with my far right, Duane Berlin is the Principal and Managing Attorney with Lev & Berlin and is the General Counsel for the Council of American Survey Research Organization.

Seated next to him is Scott Bradner, who has served in a number of roles with the Internet Engineering Task Force and is the University Technology Security Officer in the Office of Technology Security at Harvard University.

Seated next to Scott is Paula Bruening who is Staff Counsel for the Center for Democracy and Technology.

Next is Ray Everett-Church who co-authored the Internet Privacy for Dummies and Fighting Spam for Dummies and is the Managing Member of the ePrivacy Consulting.

Seated next to me on my left is Frank Gorman who is an Attorney with Bryan Cave, in the Antitrust U.S. Trade Regulation Group.

Seated next to Frank is David Kaefer, who is the Director of Business Development, Microsoft Intellectual Property and Licensing Group.

Next to him is Annalee Newitz, who is the

Electronic Frontier Foundation's Policy Analyst.

Next to Annalee is Dan Quinlan. Who is the Vice President of Apache SpamAssassin with the Apache Software Foundation.

Finally in the last seat is Jonathan Zuck, who is the President of the Association for Competitive Technology.

Thank you all for being here with us this morning. There was one change to the agenda. Howard Lipper from Morgan Stanley is not here today.

John Levine did a great job of outlining the importance of email authentication, and before we get to the technology of the different proposed standards. We must first recognize and discuss some of the policy and legal issues email authentication raises, including antitrust issues, privacy issues, and this includes the ability to engage in free, anonymous speech, and intellectual property licensing and its compatibility or incompatibility for the open source community. We're going to talk about each of these and other issues as they may come up throughout this discussion.

Let's first consider whether there are any antitrust implications with respect to an email authentication standard.

Frank Gorman, standard setting is, by its very

nature, anti-competitive, but standards are often desirable and even necessary. Here some of the proposed authentication standards are being proposed by major market players.

Now, Frank, you work in the antitrust trade regulation group at Bryan Cave, and you're also the author of Shield for Standards, which is an article about antitrust law. Can you address any of the antitrust issues you see in this scenario?

MR. GORMAN: Sure. Well, I wouldn't say that standards setting is necessarily anti-competitive but Senator Layhe put it recently commenting on the Standards Development Organization Act.

Standards development is not necessarily anti-competitive. There is, as Senator Layhe put it, unavoidable tension between the antitrust law to prohibit businesses from collusion in the development of technical standards, which require competitors to reach agreement on basic design elements.

Basically antitrust laws prohibit collusion, and standard setting requires collusion, but there are also significant pro-competitive benefits, and standard setting is now analyzed under the Rule of Reason, which means that you sort of weigh the pro-competitive, anti-competitive benefits to determine whether or not,

on balance, it is anti-competitive and therefore violates antitrust laws.

Standards are all around us. We're all able to screw light bulbs into sockets because there are standards. There are safety standards. There are thousands of standards developed on a yearly basis. They are mostly done through cooperative, non profit standard setting organizations that are essentially in the private sector.

This is essentially a government function that has been given out to the private sector, and the Standard Development Organization Act provides some protection for the standard development organizations, but not necessarily for the participants. Intra operability standards, which I think would be required in an email authentication system, can have profound positive effects on economic efficiency.

Arguably it can't work without them in email authentication. You could have a situation where you have competing models of email authentication, and then eventually what are called network externalities will come into play where there will be a typical play where one is more preferred than the other. This is what happened with Beta and VHS, if you all remember that. People who have large collections of Beta tapes recognize the

downside of that approach. That's sort of a trade market approach.

I don't know if you wanted me to get into more detail about the kinds of problems that can come up.

MS. ROBBINS: Well, standards as you put it are set all the time. Do you think that any antitrust concerns are there?

MR. GORMAN: No. People do this all the time. It's absolutely not an insurmountable problem, openness, transparency. There are problems that come up, competitors. Any time you have a standard set, you're going to have winners and losers. Certainly the example that we have going on right now with email authentication and some of the debates that are going on the MARID, it seems to be broken down into the open source camp and the licensing camp.

Ultimately, even if a compromise is reached, somebody is going to think that they got the short end of the stick, and they may bring lawsuits. They may bring litigation. There are those risks.

The different anti-competitive practices and standard setting that you see are problems with the composition of the standard setting body, improper exclusion of participants, deck stacking. The IETF has a long history of doing this, and they have good

processes in place. I think Scott can talk about that.

I did note that they have not applied, they have not filed notices with the Department of Justice and the FTC to get some protections that are available under this new act, but those protections are rather limited, and maybe Scott can address that.

Corruptions of processes is a problem. Patent ambushing where people do not reveal intellectual ownership of intellectual property can be an issue and then seek to benefit from that intellectual property, once that becomes part of the standard.

In vote stacking, there have been cases where people signed up all sorts of members for a standard setting body to get them to pass their particular version of the standard, and then the competitors sued and won and got treble damages.

Another problem that can come up, and this is probably an issue here or at least has been talked about as an issue here, is restriction of access to the standard. Some SROs can have bylaws that prevent members from owning or asserting IP rights. It's much more common to require IP rights to be licensed under what is called reasonable and nondiscriminatory terms.

If the standards are proprietary, a firm

controlling them has the power to limit or prevent competing firms from accessing that standard. They have ownership rights, and that's not necessarily a problem. What you look at under Rule of Reason analysis is whether there's a motive or intent of denial, the degree to which access to the standard is critical to effective competition, and the effect on competition from excluding the rival.

Now, different courts take different approaches to this, and some courts recently have valued the ownership of intellectual property more highly than some of the older decisions, so this is an area of law which is in flux.

If the refusal of access is not motivated to suppress competition and there's no patent ambush, the Department of Justice's and FTC's intellectual property guidelines recognize that intellectual property, like other components of production, does not necessarily confer market power and licensing rights as generally pro-competitive and efficiency enhancing because by allowing people to have ownership rights, it encourages people to do research and development.

The other problem that you may run into are standards that reduce competition by facilitating collusion or inducing incentive to compete, price

fixing, that sort of thing.

The Standard Development Organization Act incorporates OMB Circular A 119 which sets forth certain transparency, consensus based decision making, due process, sort of procedural steps that you can follow as a Standard Development Organization to be under the protections of the Act.

MS. ROBBINS: Thank you. Now, most of the proposal authentication schemes have been submitted to the IETF. And, Scott, you have served on a number of roles with the IETF, and I believe that the IETF has policies regarding the disclosure of intellectual property rights and for reasonable nondiscriminatory licenses, and do you think that those policies alleviate any of the concerns that Frank has just outlined for us?

MR. BRADNER: Well, I don't pretend to like the microphone. The IETF rules are pretty straightforward, and they don't go quite as far as you might suggest. Basically the IETF rules are you must disclose. In order to participate, you must disclose any IPR that you have, which is either patent applications or patents that you reasonably believe have to be taken into account if somebody is going to implement a particular technology, and you have to do that as soon as you know

that there's a potential problem.

You don't wait until the end. You don't wait for a last call when the standard is almost done. You have to do it immediately. We do recognize that sometimes you can't do that, and if you can't do that, then you cannot participate. You can be in the room. You cannot advocate or denigrate a particular proposal if you have not disclosed any issues that you might have, but that is pretty much the extent of the rule set.

You basically have to disclose that you have IPR or claim to have IPR, and then you have to persuade a working group that this is -- that this particular technology, in taking into account its IPR issues, is better than other competing technologies for the same application or combining it or whatever. We do not have any particular rules of what RAND means. We actually very carefully decided not to do that. We were not trying to decide whether your licenses are fair or not.

MS. ROBBINS: I would like to stop you for just one moment. Can you explain what RAND is?

MR. BRADNER: Reasonable and nondiscriminatory licensing process, requirements, i.e. everyone can license it for $10,000 a copy. That under some circumstances may be very fair and other circumstances

where it's really a patent perfect ten would kind of be hard, so we don't make any particular requirements on that, but the working group does take that into account.

We decided to avoid the question of the standards body in the IETF trying to figure out whether something was fair by dealing with our multi stage standards process. We have a three-stage standard process, that we're advising at the moment we've got this three-stage process. The first stage is a good idea, no known problems. The second stage is multiple intra operable implementation, and if there is known IPR, agreed to IPR, and by agreed to, I mean the implementers agree there are relevant IPR, not just because somebody claims there is because somebody can falsely claim, then in order to progress on the standards track, you have to have implementations which are multiple implementations that have separately exercised a licensed.

So if you both have exercised a license, then by some definition it must be fair, and if indeed there is only one exercise in license which is the one that came up with the technology, then it's probably not fair and it can't progress on the standards track.

We don't try and make a value judgment of the

particular licensing issue per se, but of course, a working group in looking at technology will take into account the capabilities of the technology, the features of it and any other factors including IPR. We do not actually require a license to be published, so that the working group will take that into account, but in general if a proposer of technology doesn't tell the working group in some level of detail what they think is going to be their licensing, it's probably not going to progress, but that's up to the working group.

It's not written in the standards process that we require it one way or another. We don't require royalty free. We don't require RAND. We just tell the working groups that they have to think about it.

MS. ROBBINS: Thank you. I would like at this time to move on to privacy concerns that the email authentication system raises.

Paula, you specialize in consumer privacy and free expression at the Center for Democracy and Technology. Do you have any specific concerns about an email authentication standard and how it will affect privacy?

MS. BRUENING: Thank you. Is that okay? Can you hear? Great. Well, first of all, I would just like to say that we think that email authentication systems

and specifically email authentication at the domain level is a really important technical development in the effort to fight spam.

CDT has long espoused the view that it's going to take a variety of different things to curb the flow of spam. One is enforcement of appropriate and effective law. The second would be the technological solutions that we're going to be hearing about over the next couple days, and it's also going to require an informed consumer and users of the Internet that there are underlying behaviors that go on that if you could avoid those, you can probably find yourself with less spam coming into your mailbox.

I think that what's important in looking at these technological solutions is to bear in mind that while this is a very important tool for commerce and we certainly recognize this, that the Internet also has -- there's been a vision for the Internet that has involved the ability of the average user to speak to a wide group of people all over the world and to engage in political speech, and sometimes that speech is anonymous political speech, and it's something we have valued in the United States for a long time.

We think that it's important as we go forward to deploy these technical solutions that we continue to

respect that ability of users to use the Internet and the email application of the Internet in that way.

However as we go forward to put these technical solutions in place CDT feels it's very important that we continue to enable people to speak anonymously when they're talking about political matters.

Now, that sort of comes out in two different ways. One is it's going to be really important that while these authentication systems are out there and they're helping to make it possible to enforce laws that require anti-spoofing, we can figure out where the email is coming from, who is doing this bad stuff online, at the same time that there is a way that people can use the email systems without having validated, authenticated email, and that that email will not be turned back out of hand.

That doesn't mean that it's the first email that's necessarily delivered. It's not that it's the quickest email, but that it's not automatically turned back and refused delivery.

Clearly the same kind of analysis is going to have to go into looking at the email to figure out, is it bulky, where is it potentially coming from, what kind of content are we talking about, but that in and of itself, simply because it's not authenticated, does not

mean it's not going to be delivered, and that's really, really important.

I think the other piece of that is that if you're going to allow this sort of anonymous political speech, there has to be an assurance that there are different kinds of technologies out there that senders can use that can really meet their own purposes and meet their own needs of delivery, whether that's reliability, cost or speed, and that there is always some kind of an open avenue for speakers on email who want to engage in this kind of speech.

MS. ROBBINS: Duane, as General Counsel for the Council of American Survey Research Organization, you deal with online privacy policies and collecting privacy information. Do you think that there is a way to balance the need for authentication -- sorry about that.

I'll start over. Duane, as General Counsel for the Council of American Survey Research Organization, you deal with online privacy policies and collecting privacy information. Do you think that there is a way to balance the need for an authentication system and balancing the need for maintaining anonymity as Paula just described?

MR. BERLIN: Yes, Colleen, thank you. I think

that actually that balancing is essential. I agree with Paula very much that anonymity in political voting and speech is important, though I think it's relevant to ask how important in relation to the other considerations we've got, and I think to do that, you've to back up a little bit and look at the way the privacy regulation has evolved in this country and in other countries.

In Europe, for example, the thrust of privacy regulation is really data protection and the ability to have control over information that's disclosed to third-parties and where that information goes.

Several years ago, when we saw the implementation of regulations like HIPAA and GLB, which dealt with the handling, use and disclosure of consumer information and how it's redisclosed and how it's used and shared, the emphasis was similar to that which we saw in Europe.

In the past couple of years, as a lot of us know, we've seen a great push in what I think is the other sort of major vein or major avenue of privacy regulation in the U.S., which is the right to be left alone. We see that of course in the Do Not Call Regulation and Statute and in the recently enacted CAN-SPAM Act, and really the subject matter of this conference, which is the right -- and that's a little

bit in quotes, the right not to receive a phone call or an email or perhaps a knock at the door or perhaps a piece of paper mail even that you haven't asked for or that you don't want or about a subject that you're not interested in.

So in email authentication, you could look at it as a very interesting nexus of those two veins of privacy regulation, that is the right to have personal data, the anonymity versus disclosure of the sender protected versus the right to be left alone or to not receive an unsolicited communication or receive information about a subject that you're not interested in or don't want to know about.

Almost by definition, almost from the get go, the subject of authentication is a balancing act between the personal information of the sender and the right of the recipient to not receive something that they don't want to receive.

It seems to me that the various factors involved in that certainly speak to authentication in the implementation of an authentication system as winning, if you will, in the balancing act between those two sets of considerations. Certainly online speech is available anonymously through other methods besides email, through the use of a web site, blogs, et cetera.

Also just in terms of the evolution of the juris prudence, the protection of personal information, that side of the consideration, that vein of the analysis, has typically been about disclosures that an individual makes to a third-party, a doctor, a bank, someone with whom they've done business and what that third party does with the information.

Typically at least in terms of the regulation that's been passed thus far, disclosures or statements made by the individual haven't received as much protection as disclosures made to third parties, not to say that that's not an important consideration.

So in summary, both sides of the equation are important. Both rights exist. No right is unknown, is exercised without some level of restraint sort of, an example being we have free speech but we don't have the right to yell "fire" in a crowded theater, so by definition I think the subject speaks to a balancing act, and I think it is soluble.

MS. ROBBINS: Paula, I think you wanted to comment on that.

MS. BRUENING: I just want to draw a distinction and make clear that what I was talking about was political speech, not commercial speech, and political speech is afforded a much higher protection by the

Supreme Court than commercial speech is, and that I think was pretty clearly borne out with the Do Not Call List where you could sign up to avoid calls from marketers, but there was a different standard for people who wanted to call you and talk to you about political matters, and I think anybody that lived in a swing state in the last couple months are well aware of the difference.

The other point I would like to make is I think there's a big difference between the power of email and the power of what you suggested in terms of blogs or chat rooms, as far as for political speech. While I agree that those kinds of tools are very important, they really don't have the kind of power that email does in terms of organizing around a very time sensitive issue.

I can't be sure that my city council person is going to come and read my blog or come and join my chat room, but I can have a better sense that they may get my email, and I can take an active step to be sure that they engage with me in some kind of political discourse in that way, so I wanted to just make those two distinctions.

MS. ROBBINS: Ray, you're the co-author of Internet Privacy for Dummies, and do you think that the domain level authentication strikes that balance that

we've been talking about as opposed to a user level authentication?

MR. EVERETT-CHURCH: I think that domain level authentication can provide sort of a level of abstraction to the authentication process that will help dissuade some of the fears about uniquely tying particular messages to particular individuals, which is a sensitive concern in the free speech and free expression issue base.

The domain level authentication does give you a much broader way of identifying the source of mail, and with that you get a level of abstraction that makes it difficult to tie a particular individual to some bad act that they performed, so there is a trade-off here, and that's why I think that it's going to require a great deal of care and consideration to apply a level of granularity that does allow a unique sender to be identified versus a domain level approach, which can give you some sense of comfort, some level of trust in the origins of the message without compromising individual privacy.

MS. ROBBINS: I just want to make two announcements. One is, if you do have a question in response to a question I asked another panelist, please just put up your table tent, and also I'm just going to

hold the audience questions until the end.

Annalee, as the Electronic Frontier Foundation's Policy Analyst, do you agree with what Ray just said, that we do need to balance the need to authenticate email and the desire to have anonymous speech?

MS. NEWITZ: No. Actually I wanted to amplify a little bit of what Paula was saying about the importance of anonymous free speech. I think when we talk about free speech and we say email is a terrific vessel for free speech, I think we tend to forget that the Supreme Court has countless times said that forcing people to identify themselves when engaging in speech, actually it has a chilling effect on that speech. In other words, having to identify yourself means that you may not, in fact, engage in important acts of speaking, political speech, whistleblowing speech.

In 1995, the Supreme Court in a case called McIntyre versus the Ohio Elections Commission said that for people to hand out campaign literature and to be forced to put their name on that literature, there was actually an ordinance in Ohio that said you had to sign your name to any campaign letters you were handing out, that that actually interfered with people's ability to engage in campaigning.

In that ruling the Supreme Court said anonymity

is a shield against the tyranny of the majority, okay? Being able to speak anonymously shields you from people punishing you for what you've said. It shields you from social approbation. It allows you to engage in the kinds of healthy acts of speaking out that are important to democracy.

At the Electronic Frontier Foundation where we deal with a lot of legal issues and policy issues and how they impact technology, we receive dozens and dozens of calls every year from people who have spoken out either through email or on discussion boards where you're identified by email and that are being actually tracked down by people who are trying to subpoena their real name, and it turns out in most of these cases it looks like that basically people are starting sort of frivolous lawsuits in order to subpoena the real names of these people just to take extra legal punitive action against them.

Let me give you a quick example. There was a case that we dealt with that was in Ohio where there's actually a law that says, you don't actually even need to bring a lawsuit against somebody in order to subpoena their real name from their ISP, which means that if you could trace them back using their email to their mail server or to the ISP that sent their mail, you could use

what's called a petition for discovery to get their true name.

So this woman, Jane Doe because she was never identified, posted in a message board that a local entrepreneur who ran this company called A.K. Steel was litigious, and he took offense at that and filed a petition for discovery to find out her real name. Now, as soon as legal action was taken to quash that, he dropped the case.

We had another case that we dealt with where a company called ToTheMark.com, which has long ago fallen off the NASDAQ, was in the midst of another lawsuit dealing with their financial situation, and so they decided to subpoena the names of anonymous speakers on a Yahoo! message board, who were just talking about how crappy the company was, and they alleged that getting the real names of these anonymous speakers would be relevant to the case, even though it turned out none of them actually worked for the company, and in fact the subpoenas were quashed. This was in Seattle.

When it was pointed out that some of these people did not work for the company, the company became a lot less interested in getting their names, and so what we find, and we found this again and again in cases like this, that basically people in these cases who are

trying to subpoena the names of these speakers based on their email addresses, getting them from their ISPs are trying to take punitive damage, usually firing them, because it's almost always people who are speaking out about a corporation's bad practices or perhaps saying that somebody is litigious who works for one of these companies and trying to exercise free speech, and they are going to suffer punishment if their real name is discovered.

That's where we come to this. We are concerned about email authentication. We worry that if people -- if the domain that sends your email is easily discovered or if it is easy to authenticate who the person is that has sent a particular email, that it will keep people from speaking out on important issues. It will basically chill the process of free speech before the free speech even begins.

MS. ROBBINS: I guess what you're saying is that domain level authentication to you has the same problems in terms of protecting anonymous speech as sender level authentication.

MS. NEWITZ: It absolutely does because all it requires -- in some states you don't even have to initiate a lawsuit, but in other states, if you do initiate a lawsuit, say a lawsuit about liable or

something, it's very easy to get the true names of those speakers, so it really doesn't provide any anonymity at all.

MS. ROBBINS: So, Annalee, looking in a crystal ball, if you look into the future and you see that the failure to adopt a domain level authentication standard results in a decrease of reliability of email, more aggressive filtering in terms of higher false positive rates and greater amount of inbox clutter that results in lost messages, do you think your answer would change?

MS. NEWITZ: No, because I think what we're talking about here, email authentication, I don't think anyone here believes that that would be the only spam solution. It's part of your complete anti-spam breakfast, right?

So what we're going to have is we're going to develop better filtering technologies. We're going to develop better bayesian filters, whatever. I'm very against commercial speech cluttering up my mailbox, just as much as everyone. Because I work on spam, I actually don't filter my mail so I can see how much spam I would get in a kind of real word experiment, so I filter through like 2,000 spams a day by hand, and it's annoying, but I don't think that the -- yes, I suffer

for spam.

But I still don't think the collateral damage to anonymous free speech is worth it. I think what we need to do is focus on other kinds of technology that will stop spam.

MR. GORMAN: Annalee, I think you're making some really strong policy argument, but I wonder how you get around the State Action issue when you say that it violates constitutional free speech to have some sort of domain level authentication. I don't see any State Action there as long as it's done by the Standard Development Organization and not by the government.

Again I think you're making very good policy arguments, and I think they need to be taken into account, but I don't know that it rises to the level of constitutional violation.

MS. NEWITZ: I think it's going to depend on the context. I think that in some cases, you're absolutely right, and I think it is -- I really do want to make this as a policy argument. I'm not claiming that if we institute email authentication, there's going to be this reign of sort of Constitutional violation problems, but in some cases I think it is possible that one could argue this is violating First Amendment so I think that's a huge risk.

MS. ROBBINS: I'm sorry. Dan, you have a response?

MR. QUINLAN: I just had a question more so for Annalee. So one thing I guess I'm confused about is that you say that authentication would make the problem worse than it is today, but people already today are subpoenaing domains. I guess I'm confused about how authentication would change the landscape as it is today in terms of reducing the possibility of anonymous speech, and it seems to me that anonymous speech is still very possible with a domain based authentication.

There's no need to tie some authenticated entity with a particular individual. As long as an authentication scheme preserved that ability, would that alleviate some of your concerns with it?

MS. NEWITZ: It might alleviate some of my concerns, but let me answer your first question first, which was would it make it worse, and I think, yeah, it would because what we're hoping for is a situation where pretty much everybody is engaging in some kind of authentication because that's how it's going to work best.

If that's true, that means every email sent can be traced back to its domain of origin, which is a different situation from what we have now, and I think

it would make it easier for people to subpoena those true names if they always know what domain this email is coming from, so I think that's a danger.

Your other point, if you're just tracing it back to a domain but not to a particular user, again if I'm say Annalee@example.com, but I also go by Biffy@example.com and Scoopy and Whippy@example.com, you're still going to be able to trace me back to example.com, and if you subpoena them and you say, who is Annalee and Scoopy and Whippy and all those other names, it's likely that they are going to have some kind of record that traces it back to Annalee Newitz, so that's my concern.

MR. QUINLAN: I mean, even today you can identify exact IP address that a message came from. It seems like that's even easier to track down than a domain, and authentication schemes are not going to make that became unavailable.

MS. NEWITZ: The kinds of people who are trying to subpoena these names are not necessarily the kind of people who even know what an IP address is, so you're talking about people who are like trolling on a Yahoo! board or who are on an email list, and they see a mail, and they say, well, I don't like what this person is saying about my company on this mailing list, I want to

find out who they are, and I know that they come from example.com because SPF tells me.

So I go to example.com with my subpoena, and I say, I'm bringing a suit alleging defamation of character and I want the name of this John Doe who said that my company stinks, because they're hurting my business and they're potentially lowering my stock price and give me their name.

So that's sort of the nature of my concern. I don't know if that answers your question or not.

MS. ROBBINS: Annalee, we now have sender level authentication for our telephones. Do you think that email then should be treated differently than our telephone systems?

MS. NEWITZ: Well, we don't force everyone who makes a phone call to identify who they are. We have Sender ID on phones, but you can turn it off. You can also spoof it and that's legal, so far.

MS. ROBBINS: Ray, you had a response you wanted to give?

MR. EVERETT-CHURCH: Yeah, I just wanted to say that I agree with the speaker a moment ago who complimented the policy arguments. I think they're a very important part of this discussion.

I just wanted to add that I think, I may be

incorrect, but in my review of all of the major authentication proposals out there, I think even the most rigid and robust authentication schemes being discussed have, within their frame work, the capability to prevent remailing systems, to permit certification programs by third parties who might be able to vouch for someone who is seeking to communicate in an anonymous fashion an EEF or an ACLU or some other entity who can provide an umbrella domain that may, as Paula said, not get on the fast track for delivery of a piece of email but wouldn't necessarily be caught in a vacuum there created by the need to have some sort of domain level authentication.

So I may be incorrect, but I think most if not all of the proposals have within them the capability to permit anonymous communication, and if there are proposals that don't, I think they should be severely questioned because of it.

MS. ROBBINS: Duane, do you have a comment?

MR. BERLIN: I just wanted to pick up on the comment that was made before about the implication of State Action and how that's necessary in order to generate a First Amendment issue, and I think that that corresponds to what I think is implicit in the discussion, and that is that most of us believe that

commercial email that's abusive and the source of which is concealed is objectionable and should be regulated and that our main concern is about personal and political speech.

As was mentioned in the opening comment, the teeth of an authentication system occur when a regulation is implemented that would make it illegal to hack into the authentication system. If that regulation speaks to commercial email as the CAN-SPAM Act does, as Do Not Call does with respect to commercial phone calls, as the Telephone Sales Act and the TCP Act do, if the State Action speaks to commercial email that is deceptive because the center is concealing their identity, then I think that goes a long way to beginning to make the distinction between personal political speech and commercial speech that we're sort of wrestling with here.

MS. ROBBINS: Ray, in terms of the effect on anonymous speech, do you think it matters whether the authentication standard is IP based or signature-based?

MR. EVERETT-CHURCH: Again I think that the most important consideration is that whether you're considering an IP based solution or some sort of digital signature approach, that you have within that framework the capability to support anonymous speech and free

expression.

You've got to keep these considerations in mind as you develop these proposals and as they move forward through the standards process, and it's something that I think that the industry also needs to bear in mind because I think there may yet be some business opportunities here for tools that will enable entities to act as an agent for those who are seeking a reliable way of speaking individually and potentially anonymously.

There are tools that could be built, designed, whether this is an IP approach or rapid approach, that would give end users some better ability to control how that mail comes to them, how it flows through, filters and blocking, et cetera, to ensure that they do get the types of communications that they're seeking and that those communications aren't inadvertently impeded because of a problem meeting an authentication standard.

Certainly I think IP level approaches have some of the broad capabilities or broad features of a domain level approach. There's some bit of abstraction there, but then again digital signatures can be signed for an individual or for an organization or for a range of organizations. There's a lot of granulatory there.

MS. ROBBINS: I'm going to switch gears now and

focus on some of the legal issues dealing with patent licenses. There are at least two patent licenses available for authentication technology. Yahoo! has a patent license available for DomainKeys, and Microsoft has one available for Sender ID.

There have been issues raised with respect to the software patent licenses and their compatibility or incompatibility with open source software, and I would like to take some time now to discuss this issue further.

David, you are the director in Microsoft's IP and Licensing Group, and Microsoft is offering a patent license for when or if a patent is granted on one specific portion of Sender ID, the purported responsible address check. Could you explain why in that license Microsoft includes reciprocity and defensive rights provisions?

MR. KAEFER: Sure. I'm happy to do so. It would might be helpful to start at a little bit higher level as well to give you some insight into our overall framework as we wade into the IETF process, as well as the framework that other companies have as they go in and make IP contributions in similar forums.

You know, we have 25 years or more of IP standardization experience in the software industry and

25 years of success in dealing with patent issues as they relate to the standard setting process, so there's an awful lot of norms and standards that people can look to over a period of time to sort of determine what is common within a license.

As we went about and crafted the license that Microsoft is providing for its patent application that is relevant, as Colleen mentioned, to one segment of Sender ID, first what we looked to was really what are the norms, and we tried to stick close to those.

Certainly standards licensing has a number of different things to look at. One is, as others have mentioned, the royalty basis. Is there going to be a charge associated with somebody contributing a piece of patented technology? The answer is it just depends.

Some people decide to contribute their IP on a royalty bearing basis. In this particular case Microsoft, in the interest of making sure that as many people as possible can use its patent application, wanted to make sure this is available on a royalty free basis, so certainly royalty bearing issues are one thing to look at.

Reciprocity, Colleen, as you mentioned, is also a very basic principle that is in most standards related licensing, not just in the existing IETF work that we're

looking at today, but pretty much in all. By reciprocity, really what we're talking about is everybody who is participating in the standard agrees essentially to provide similar rights back to people who are contributing IEP to the standard.

So, for example, if party A contributes a right on royalty free grounds, other parties who want to actually use that right would essentially provide any necessary patent claims that they may have with respect to the patent or a patent application back on similar terms. That's very important because everybody should be playing essentially by the same rules, and essentially that's what reciprocity does.

The positive affect of reciprocity also in the standard setting context is it sets up a legal framework, if you will, for people to do business with one another, for people not to end up in a situation where there are legal disputes because it encourages all people in this case to contribute to IP in similar ways and understand by all folks who are participating.

When reciprocity breaks down, when there are bad actors and there's a patent holder who either in the context of the standards participation or a standards holder who is not participating in the standards body decides to come in and litigate against anyone frankly

who is implementing standard, whether it be somebody like in Microsoft's case is contributing IP or frankly just somebody else who is implementing in this case a Sender ID spec, and that's a bad outcome.

Reciprocity helps essentially reduce the likelihood of that type of dispute.

MS. ROBBINS: Can you also explain or give an example of what would happen if you didn't include those provisions within your license?

MR. KAEFER: Again I think the central point here is that all people have to play by a set of common rules, and the only way to make sure that everyone is playing by the common rules is that everybody participates actively in the licensing of that IP.

One issue that's come up within the context of this particular IP license provided by Microsoft is this notion on sub-licensing, which is actually one of the central questions with respect to some open source implementers.

Now, sub-licensing essentially is this concept that if A provides a piece of IP, in this case a patent application through the standards process, and B decides to license it and implement a spec on it, that the sub-licensing prohibits B from passing that license forward to another party, party C.

Now, why is that important? Well, we don't know who C is. C is at arms length. C hasn't necessarily negotiated an agreement with A. We don't know what rule C is playing by. We don't know whether or not C has decided, for example, to contribute its own IP on a royalty free basis but in similar terms, in a reasonable nondiscriminatory way adopted by the standards organization.

By essentially encouraging everybody to participate in that process, you're bringing everybody in under sort of a predictable legal environment.

MS. ROBBINS: Jonathan, you are a professional software developer and also president of ACT, Association for Competitive Technology. Could these provisions that David just outlined be seen as a benefit to the licensee as well as to the licensor?

MR. ZUCK: Thank you, and thanks for the opportunity to participate today. I mean, as David mentioned, IP has danced well with standards process for a very long time with a great deal of success, and I think it's always important to take a step back from a theoretical discussion and have a practical discussion about these issues, and one of the key components of some of these provisions is kind of an inoculative effect that you provide.

When you have a situation where reciprocity is

the environment of a standard, then you're less likely to have a more litigious kind of Johnnie Come Lately patent dispute because you've created a community of people who have all agreed to contribute their IP on reasonable and nondiscriminatory terms, so that kind of environment is actually beneficial to everyone involved in implementing the standard, not just someone providing a specific piece of intellectual property.

So, the practical implications, there's nothing about these licenses that represent true barriers to adoption of the standard, and the protected benefits far outweigh any of the inconvenience that might be associated with downloading a license, signing it and faxing it to a company that's contributed IP.

MS. ROBBINS: Scott, I believe you wanted to comment?

MR. BRADNER: Yes, I would like to back up a little bit and talk a little bit about what happened in the IETF relative to these licenses that were spoken of.

The IETF had a working group which was working on thinking about Sender ID and similar technologies, and Microsoft provided an intellectual property right disclosure and license, which actually exceeds the IETF's process requirements. There's no requirement in

the IETF process to provide a license, but Microsoft did.

And the license was, as you've heard, for royalty free with reciprocity and no sub-licensing and actually executing a physical license. It was an unusual license relative to the IETF because we had not had one before which had the no sub-license or the executed license be required, but it's not irrational in the sense that it doesn't violate any rule set.

The other thing that Microsoft did in providing the license was they provided a remarkable tool for confusion. The license was written in lawyer. It wasn't written in human. The geeks that come to the IETF just didn't vaguely understand what this license asked for.

I participated as the author of the IETF's intellectual property right rules and processes. I got involved in this at the request of the Chair and at the request of other people involved. I had one exchange with a system manager at a university who said that, well, under Microsoft's license, the university would have to give up its entire patent portfolio, including biotech patents, in order to run this software, not to modify and distribute it but even to run it.

The license was extremely difficult to read for

non lawyer types, and I think that 95 percent or more of the discussion over these licenses was completely not a reality. It had to do with misunderstandings of what the license was asking for, so Microsoft did itself a disservice in providing that license because of the way it was written.

It went beyond the requirements of the IETF in providing licenses, but the two provisions that caused the most difficulty, specifically in the provisions of having to execute a physical license and no sub-licensing were seen by parts of the community, the open source part of the community as unacceptable, but not all of the open source community felt that way, but enough of it did that this was a significant issue.

The MARID working group was closed but that was not the reason. The MARID working group is looking at multiple technologies to work on a particular part of the anti-spam problem, and there were significant technical disagreements over the specific technical proposals independent of the licensing issue, and it became clear that the working group was not going to reach consensus on the technology itself independent of the licensing, and so the working group was closed.

Notice that in the IETF, working groups come and they go. They're not standing committees. It's not a

big deal to close a working group, so it shouldn't be taken as some cataclysmic event because we do it all the time.

We have in the past had a number of cases where we've failed to make progress in an individual working group when there are competing technologies and groups within that working group that aren't going to compromise within that, and we found that in creating multiple working groups and proceeding on that basis has been much more successful, and this may happen in this case.

I want to be sure that people understand that it wasn't closed simply because we got into a food fight over IPR. It was technology as well, but I do want to say if you are writing a license that's going to be going in front of a standards body, please write it in something that the geeks will understand. Whatever this was written in, it used English words, but not in sequences I've run across before.

MS. ROBBINS: Thank you for explaining that. Actually, Scott, I have a follow-up question for you. If the license was understandable and reading it as it is now, do you see a patent license that contains the terms that this one does as necessarily a bar to adoption?

MR. BRADNER: Again I would like to back up one little bit first, which is the IETF does a lot of work, a lot of standards which have IPR disclosures and claims on them, and there are many environments where RAND as in not royalty free but actual licensing terms is just fine. We have a number of technologies where every single proposal made to the working group was something that somebody wanted money for, and the working group looked through it and worked out the best set of technology they felt could do the job and then proceeded with standardization of that, even though there's royalties that are going to have to be paid.

These are technologies, for example, that cell phone manufacturers use to make cell phones, and they know about this anyway.

There's another category of the technology that IETF works on and that is so the core infrastructure technology, TCP itself, the web, emails, things like that, which a great deal of that technology is implemented in open source. It's not implemented -- it's not merely implemented in large commercial companies that sell the software, but it's by open source, Apache and the web domain, Sendmail and the email domain are major players in this.

So looking at licensing terms and licensing

characteristics in those two different areas are very different, and it's not easy to characterize the IETF as being royalty free or whatever simply because we cover such a wide territory.

In the face of the kind of thing we're talking about here which is something that is the implementation of which is going to be dominated by a mixture of open source and commercial, we have to take into account the open source. As I mentioned earlier, not all of the open source community found this particular license to be impossible to deal with, but some of it did.

Some of that probably came from a generic distrust of the open source community, Microsoft for reasons I don't need to go into, I suspect. I don't know. I'm not a lawyer for the open source community, but some of the lawyers for the open source community said that the non sub-license was simply not something that they could deal with.

The license itself, having to execute a license, is probably something that most of them could deal with. At least ones that I talked to said they could, but they said they could simply not deal with this non sub-licensing, but there you have to talk to the people who actually are saying that, who are actually in the community, and the ones that talked to me said it was

not possible.

MS. ROBBINS: Dan, I have a follow-up for you about the sub-licensing. Do you want to respond to that first?

MR. QUINLAN: A couple things. First to go back to the IETF processes and the reason that the MARID working group closed, I would say that it was actually the case that the primary reason the work group closed was over the patent license for Sender ID and the portions of Sender ID contributed by Microsoft.

If you look at other working groups that have closed, and there are groups within the working group that have a fundamental disagreement that does not get resolved, it's typically over the standard technology itself, and in this case there was more of a rough consensus around the standard technology, but the complete lack of consensus that I think and I think most people would agree caused the work chairs to agree that the working group needed to be closed because it wasn't going to succeed in producing a standard which was the patent license.

Regarding the assertion that the main problem with the license was that geeks cannot understand the license, luckily the Apache Software Foundation, the organization I'm here representing, worked with an

attorney who was able to understand licenses, and we do have a long history of open source licenses, and these are not unknown things to us geeks here.

So I would say the majority of the people on the mailing list who disagree with the license actually really understood what the effect would be on their software and how it affects the competitiveness of open source software in the marketplace, and the main objections were because people were concerned that this would suppress open source software and make it more difficult to distribute their own software.

In terms of how wide a group was concerned about the license and making sure the email authentication was available, it was not just Apache Software Foundation but also the Open Source Initiative, the Free Software Foundation and Software in the Public Interest, which I think are probably the four most significant, open source nonprofit establishments out there right now, so with that cast of the characters saying that there are concerns about the licenses, it seems pretty hard to dismiss it as just 5 percent or a few people disagreed.

And there were people that -- there were open source developers that thought the license was a problem, but I'm not actually certain what open source software they represent or what programs they had out

there in the Internet.

So I think just to step back a level and talk about what our primary concerns are in email authentication, our main concern is we want to see an open and competitive landscape for authentication standards. Distributed systems such as the Internet are very good at picking technologies such as email and the web.

If you look at the history of the Internet standards, there are technologies such as Gofer, which maybe not everybody here remembers, but there was a brief moment in the Internet where Gofer was the way you browsed the Internet and navigated, very similar to the web, no pictures, and when pictures were available, that took over, and that's the world wide web that we have today.

The Internet made that decision on its own. It didn't require royalty free patent licenses or anything. Those standards were available for free with no licensing terms whatsoever, and the distributed system made that choice.

If you look at the Internet today, this is an example of what type of competitive landscape we do have and how open source has been successful, the Apache web server is now run by the majority of the web servers on

the Internet, and that is possible because the world wide web and the standards that are needed on the world wide web are freely available.

There's no patent license that needs to be executed with Microsoft or any other company, and we want to make sure that it stays that way for email and other important parts of the Internet.

MS. ROBBINS: Before I get to -- I have several presenters that want to make comments. I want to ask you, Dan, if you can briefly explain why non sub-licensing is so important to the open source community.

MR. QUINLAN: The main issue of sub-licensing is that the refusal to allow sub-licensing in a standard that needs to be implemented in open source software that forms the core of the Internet infrastructure is that allowing sub-licensing reduces friction for open source.

If you inserted requirements for each distributor to execute a license separately and that would basically get in the way of success of past open source efforts that have led to problems such as the Apache web server, SpamAssasin, it would be analogous to, for example, if you look at -- I don't mean to pick on Microsoft, but they're here at the table, Microsoft's products, they provide a wide variety of

open source products in their own products, and I believe they continue to do that.

And if they were required, for example, every time somebody wanted to distribute their software or sell it into the store, that the person that was distributing it needed to sign an agreement with BSD or the Free Software Foundation, another organization, I have a feeling they would not be in favor of that, every time you wanted to open a store and sell one of their products, that somebody would have to execute an agreement.

So reducing that friction is really needed for open source software to compete in the landscape.

MS. ROBBINS: Scott, I believe you were the first one to have your table tent up.

MR. BRADNER: I think that you and I read different mailing lists. I don't think that the geeks understood the license, but I'm going on why the working group closed from a direct conversation with the area director that closed the working group yesterday, and I can't be in his mind to be sure he was telling me the truth, but he was extremely clear that while the IPR was an issue, it wasn't a reason.

MS. ROBBINS: And, Jonathan, you had a question?

MR. ZUCK: Well, first, I think we can all agree that Apache has accomplished a lot of incredible things, and I think the question I would turn back to Daniel eventually is exactly how a license like this would have prevented Apache having the success that its had.

Again it's very easy to raise the kind of theoretical objection to a patent license, and I think it's interesting that he's talking about geeks understanding the license and then started talking about all distributors not able to distribute the software when in reality that's something that's explicitly allowed in the license he's talking about.

This license is basically saying if you're a new implementer of that technology, not just a distributor or indirect distributor, new implementer, somebody that's putting out their own product, that they're required to execute that license, and that's exactly the context in which the reciprocity would be so important. It's not about some store distributing it. It's about a new implementer of that technology.

Again IP has been an integral part of the standard process for a long time, and that's including the open source community, and the open source community has managed to thrive in an environment that coexists with IP. Most major open source package vendors sell

specifically software that isn't covered under the GPL, for example, that goes along side the software.

It finds a way, vendors find a way. There's absolutely nothing, nothing in this license that would have prevented Apache to have the success that it's had today or SpamAssassin to have the success it's had today, and it's important to get specific and practical about this because of the severity of the spam issue that we're all trying to confront.

This is just a first step. This is just the beginning of what we need to do to start to combat the spam and phishing problem that we're here to discuss, and there isn't a valid barrier to adoption, it's easy to adopt. It's very few people that would need to be signing a license, only people that are producing their own implementation of their own software development.

MS. ROBBINS: David, you wanted to respond?

MR. KAEFER: Yes, one I think it sort of bears some time to talk about the collaboration that took place at IETF both with Microsoft and with other commercial vendors as well as various members of the open source community.

I think it's important to note that everybody at the table recognizes a couple things. One is that the open source community is here to stay, and they've been

very successful doing a lot of very good of good things.

The second thing, a lot of people recognize that IP not just an inconvenience to be ignored. Patents in particular are something that you have to deal with head on and you have to deal with as a real issue, and there are particular ways that the industry for a long time has dealt with those issues.

Now historically the open source community has not participated in some of the more patent heavy discussions that the industry has had, but increasingly, both for Sender ID and other kinds of circumstances, we're starting to see patent issues and open source issues coming together, and there's going to be some roadblocks for folks to try to overcome.

The reality is a lot of open source licenses were created at a time when open source was not utilized in commercial settings. As open source commercializes more and it wants to use more and more patented technology, there's commercial realities that come along with that.

Now, with respect to people who originally crafted some open source license and the general public licensing being among them, one of the chief objectives of crafting that license was essentially to create a patent free zone within the general public license

source code community.

Now, that's not to say that that won't change in the future. That's not to say that collaboration can't happen, but there's over time been a desire to keep software patents and open source separate. To the degree that we all need to work together, we have to find some kind of common ground in which to make that work happen.

So the good news is it's been a good working group that actually provided one forum for us to work with some of these issue, and particularly we valued a lot of feedback and collaboration with Sendmail, for example, who I think participated in a very positive and collaborative way throughout the process, and we worked with them.

We said, what are your concerns, and I think we addressed some of those concerns, and some of the concerns we weren't able to address, and that's sort of the normal course of any negotiation between multiple parties. We're going to be able to bridge the gap in many things, and some things we won't be able to bridge the gap.

It's certainly the desire I felt on the open source community side as well as the desire on our side is to try to reach common ground if possible, and I

think despite the fact that we might be focusing today on a few of the areas where we disagreed, the important thing is to recognize the common desire by both sets of interests to work together.

Now, with respect to one of the points that Dan brought up, I wanted to clarify a couple of things. One is the Microsoft license explicitly allows end users and the people who are simply distributing trademark licensed product, it does not require them to sign a separate license. The license is very explicit about that.

So with respect to the example you provided, for example, on what Microsoft might be comfortable doing is it provides its products through our channel partners and then on to end users. That's not really an example that I think fits given the terms of the license.

The other thing that I think is important to recognize is one of the explicit points of feedback that we certainly heard from the open source community was the desire for us not to place any restrictions for folks who wanted to implement all the open source license rights that they feel are important, the right to see source code, the right to modify it, the right to redistribute it, and in fact many open source licenses explicitly require that there not be additional

licensing requirements passed forwarded either to the immediate party that takes a license or pass forward to sub-licensed parties as well.

This is something that frankly I think was the result of some of our collaboration with the open source community, but I want to read a part of our license for you, to make absolutely clear that we're not placing any obligations on Apache or Sendmail or anybody else in the open source community to take this license from Microsoft.

The core point in our license is this: "For clarification, this agreement does not impose any obligation on you to require the recipients of your source code implementation, of license implementations to accept this or any other agreement with Microsoft."

If you would take a look at some other licenses that have been forwarded by Yahoo! and forwarded by other companies, they take a different approach. They actually require you to pass forward some of these requirements on to your sub-licensees, but we understand this is something supported in the community, and I think it's something we can work collaboratively together to address.

So as I look at it today, what I see is a lot of open source licenses that will work very well with the

license provided by Microsoft, the BSD license, I think the Apache license, though I understand you've made some changes recently, the IBM Common Public License, the MIT license. All these are licenses which certainly we believe work and given the flexibility the open source community has shown on licensing over years, the fact that there's over 50 approved open source licenses, there's certainly a great deal of chance both in the open source community and within the standards context to find something that will work for everybody so long as there's a willingness to find a solution.

MS. ROBBINS: Dan, you wanted to respond?

MR. QUINLAN: Yeah. We were very willing to find a solution. We worked with Larry Rosen, who is the General Counsel for the Open Source Initiative, and negotiated with Michelle Herman, an attorney at Microsoft for several months.

Before we bog down the two major issues of sub-licensing and separate execution, the claim that separate execution is not a problem for open source ignores one of the fundamental reasons that has been successful is we don't distinguish between types of use and types of users. All users of open source are granted the same rights. There's no end user versus distributor versus someone making changes.

They're all given the same rights and not required to execute additional licenses on top of our license, so while it's fine to say that if we send the Sender ID license, the patent license, that we would not have to require our distributors to sign a license. In effect they are still required to get a license from you if they are infringing on the patents that you're claiming, so unless they're an end user since you distinguish between end users and distributors.

I think it's important to go back to comments someone made a little bit earlier which is talking about the norms of Internet standards, and why I think that MARID was actually a success and the IPR process actually worked in a way, because most Internet standards are especially for core infrastructure that if you open the open source work, that there be a competitive landscape in the field.

And in this case the IETF worked because when there was a potential for a non reasonable license to get adopted by the IETF, they shut it down, and it didn't happen, so I think the IETF process actually worked quite well in this instance.

MS. ROBBINS: Jonathan, you wanted to say something?

MR. ZUCK: Yes, and I don't want to beat a dead

horse, but the W3C is another organization that's become very eminent in the Internet space, recently went through a huge negotiation over IP practices. Larry Rosen was part of those discussions and at that time had no difficulty with reciprocity or sub-licensing provisions as part of the IP rights negotiations in the standards process.

Again I think it's important to separate the theoretical from the practical. Yes, theoretically every user of open source is a distributor. Is that practically the case? No. We know the practical realities are that there's a definite minority of open source users in fact become reimplementers or redesigners and redistributors of software.

It's that practical reality I think we need to remain focused on in the context of finding this compromise between Microsoft's legitimate or any other company's legitimate desire to protect their intellectual property and to preserve defensive rights in the context of litigation.

Let's not forget that the extent to which Microsoft preserves it's defensive rights, it created a less litigious environment for the open source community as well. The other people that might want to assert their IP rights late in the game that have accepted this

license with reciprocity are more limited in their ability to sue not only Microsoft but everyone else that's an implementer of the standard.

Again we have to separate the religion from the practicality of getting the spam problem solved and getting started down this road, and I think no one has really been able to point to the practical barriers to adoption of the Sender ID standard, and certainly the publication of SPF records, which everybody should be doing now anyway, isn't even in question.

I think it's also important to remember that everybody today can publish SPF records. Everybody can check Mail From as a means of authentication, and that there won't be any discrimination against email if you publish those SPF records. This license is just about one particular way of authenticating email, not about how you sent it, and that's really the function of whether or not any discrimination would occur out in the email space.

So again practically speaking, it's an easy standard to implement and one that I think we should get going and doing.

MS. ROBBINS: John, I guess what you're saying is that the purported responsible address check is covered by the license, but the Sender Permitted From is

not covered by the license and that implementers of Sender ID could choose to check only the SPF and not choose to take a license; is that right?

MR. ZUCK: That's exactly right. There can be plenty of debate about whether PRA, is superior and whether other technologies are coming down the road will be better still, but the foundation of this is the publication of the SPF records in the first place that will in fact be the records that everyone will be using to check whatever means they may check, and that doesn't require a license by anyone, and that's the thing we ought to start doing today to get started down this road of authentication.

MS. ROBBINS: I think, David, you had a comment you wanted to make first.

MR. KAEFER: Yeah, I just wanted to clarify one other thing as well. Let's be honest, IP licensing is not something all of us wake up in the morning and think, whoa, what an exciting topic, I want to drill down into this, but it is nevertheless a very complicated one and one that is prone to misunderstandings and prone to all sorts of different things that you learn over time.

One of the things you learn in standard setting is that even when a license is made available,

frequently a lot of people just don't choose to take the license and they just remain sort of in a limbo state, which is to say that they're not necessarily licensed for a specific patent but they're not unable to obtain that license either.

That is frankly by and large how a lot of standard setting happens. A license is made available. Some parties, because they want to get certainty around their rights, will go ahead and take a license others may choose not to. That's always an option.

In this particular case of course we want as many people as possible to agree to a license because obviously that does reduce the likelihood of legal disputes in the future, certain bad actors, either inside or outside the standard space, but there's always the option to move forward without that.

And the license that we make available today is the license that we're going to make available to anyone at any time on into the future as well. So if you don't choose to take a license today, maybe you'll choose it in five years or ten years or whenever you feel like you might need to do so.

The only thing I would sort of finish on because I think the horse is fairly alive, but I see one leg kicking, I just wanted to sort of underscore one of the

things that Jonathan talked about which is this notion that you have to find real world solutions that work for the broadest set of people possible and you try to make that happen as best you can. We're here today to solve a very perplexing problem. It's our customer's number 1 problem, which is the email is not very productive today for them because so much of it is unwanted.

We have a technology solution. The technology solution in Sender ID is something broadly, both AOL, Earthlink, Microsoft, Sendmail and others all have expressed a willingness to go forward and adopt and utilize. We have technology choice within what we're talking about, and that technology choice also allows us to steer clear of some of the their error IP disputes, which unfortunately we've had to discuss and is productive to discuss today.

Nevertheless there are ways around that, and I think what's important is to realize we have a practical solution that's ready to go that can be implemented today. We can have a real world positive impact on customers, and one thing I did want to make sure we don't lose sight of the fact that this is about consumers at the end of their day and their best interests.

MS. ROBBINS: Before I get to your comment,

Scott, I just want to ask Dan a question. If Sender ID does emerge as the email authentication standard with the licensing intact, do you think there will be in effect on the open source community's ability to compete in the email space?

MR. QUINLAN: I think it may have a negative effect. I can't say for certain that it would, and I would encourage people to explore SPF and to publish records for it to see how well it works. SpamAssassin currently supports SPF, and we do SPF checks based upon the unincumbered portion of the Mail From.

It is kind of a concern to us that Microsoft has said that they will not be fully supporting the Mail From portion of the specification and will be encouraging their vendors and partners to only support PRA fully and incumbered portions of the spec and to not fully support Mail From, although they are encouraging people to publish records, which is good, but it does kind of seem that they're saying there isn't an issue, and open source community has nothing to fear, but we want people to only really fully support the encumbered part of the spec, and given some of Microsoft past statements about open source, I think it is reasonable for us to be kind of concerned about that.

To talk for a moment about some comments that

Jonathan made, reciprocity is not one of the major concerns that we have with the licensing. If you look at our new Apache license, the new version of it, it does have some similar defensive claims around patents and technology contributed to Apache, so that is not one of our major concerns. We're more concerned with the sub-licensing and the separate execute requirement.

MS. ROBBINS: Don't those provisions though help in terms of the defensive right so that you can't sue someone unless you have them signing an executed license?

MR. QUINLAN: That is the position that Microsoft has taken. Our attorney disagrees with that essentially.

MR. KAEFER: I've never heard of that happening before, attorneys disagreeing.

MR. QUINLAN: One other real minor comment about the W3C, we actually are or probably me more personally experiencing because I'm not sure what the Apache position is on this, but the W3C patent policy is excellent, and if it included sub-licensing, then it would be perfect.

MS. ROBBINS: I know, Scott, you wanted to make a comment.

MR. BRADNER: Just a couple little things. One

thing, I thought it might be useful to know, we've been focusing on a particular license being offered and an IPR statement being offered by Microsoft. It might be interesting to note that within a week or two when Microsoft made that particular statement about licensing, Cisco also provided an IPR statement about a core technology, a way to secure TCP itself, and they took a somewhat different approach, and I thought it would be useful to just show that kind of different ways you can do things.

Cisco's approach was if indeed these standards were adopted, then anybody could implement it under RAND and went on to say, but we define RAND as being, we will not enforce the patent against anybody who doesn't sue us, and that specifically means an open source -- as long as open source doesn't decide to sue Cisco over implementation of an IETF protocol, then anybody can use it, and Cisco simply will not enforce it.

That's a different take on it, but even that take, just to set the stage of the sensitivity to IPR, in standards processes including the IPR, even that took a great deal of discussion in the working group to get people to understand what the implications were and what the issues were on it.

In the end, the working group offhand decided

that it was reasonable enough to continue to work on this technology, despite the -- again it's sort of a patent application on a patent, so I think that was just an alternate way to approach the same problem.

MS. ROBBINS: Scott, I have a question for you. If Sender ID's license or license terms stay the same with the non sub-licensable provision, is it possible that Sender ID will be adopted on a scale large enough to be effective?

MR. BRADNER: I couldn't tell. That's an open source issue. As I said earlier, that we have relatively few players in the software business for this category of core function, some of which are commercial and some of which are open source.

If some part of the open source community at the end of the day believes that they can cannot implement it because of some of the provisions of some of the licenses it certainly will do things. It will not likely effect the vast majority of users which are on -- email users who are on Hotmail and things like that.

It's more of a question in the enterprise space. A lot of the enterprise space is using commercial product of one kind or another, even if it's a repackaged open source version.

So I can't answer that. I don't think anybody

can answer that. I think it's a theoretically worry. Whether it's a practical worry or not, you can't tell until it happens.

MS. ROBBINS: Okay. I see that the time now is 10:30, so we only have about 15 more minutes left on this panel, so I would like to now open it up to audience questions. The gentleman in the red tie?

MR. MCCARTHY: Is this on? My name is Mark McCarthy. I'm from VISA. I would like to thank NIST and the FTC for this workshop, and, Colleen, you in particular for this panel. I've been in a lot of FTC workshops, both as a participant and in the audience, and I think that this was an excellent one. It really did focus a lot of the issues that we need to focus on in order to go forward in a cooperative fashion.

For us at VISA and with other financial institutions, I talked to some of my friends from American Express before the panel, we all have concerns about getting this issue moving forward. For us the problem is spam, of course, but phishing is a major concern for us.

We've taken some steps to protect ourselves. We worked with the FTC and the Treasury Department, the Better Business Bureau call for action back in June to get consumer education out there to let them know about

the dangers of phishing emails and the frauds involved, but it's even better if email authentication at the domain level can make sure or at least make it less likely that these kind of emails don't get into consumer's inboxes to begin with, so we think it's important to move this process forward.

My questions are two: Number 1, in terms of standard settings. We've all heard the discussions about IP and other concerns, and VISA is one of the biggest IP holders, patent holders in the financial services world, so we know those issues, but what's the process for moving forward in terms of standards? Is it -- maybe Scott, is there a reopening of the IETF working group or multiple working groups?

The second question is on the issue that was talked about about an hour ago, so you've probably forgotten about it, but it's the political speech question. There the issue seems to be if you can figure out the sender of the email, then it's only one small step to finding out who the actual person is.

This is to you, Annalee, and to Paula, is there a role for the federal government to set standards for the terms and conditions for access to that kind of information once you found out the identity of the email sender?

Thanks very much.

MS. ROBBINS: Maybe, Scott, do you want to take the first question?

MR. BRADNER: The people in the IETF have not stopped thinking about this question just because the MARID working group was closed. There are other activities. We are going to be involved in another aspect of that at this time, but it's been delayed until the next IETF meeting.

I fully actually expect more work to come forward, and as Dave Crocker, who you're going to hear from later today and I think tomorrow, has put it: That the IETF is good at taking something where we understand the problem and understand the set of solutions and working out the details of the solutions, no standards body is particularly good at inventing new solutions on the fly.

There are other solutions for different parts of this problem, which are coming and re-gelling, and as they do gel, the IETF certainly is going to be pursuing those areas and standardizing in those phases, once we understand them better.

MS. ROBBINS: Paula or Annalee, do you want to address the second question?

MS. NEWITZ: I can. There are already laws that

govern how people can gain access to the true names of individuals that have sent out any anonymous email. It depends on your jurisdiction, but generally there needs to be some kind of lawsuit that's been initiated, and in most of the cases that we see, it's almost always some kind of libel or defamation or trade secret type suit because these are usually whistleblowers or people complaining about companies.

So in that lawsuit a subpoena must be issued. It's a subpoena for subscriber information, which is any information you've given to your ISP or whatever group it is that's managing your email at the time that you subscribe, so that could be as little as a name. It could be as much as name, address, phone number, all kinds of other stuff.

Generally the practice now is that when your ISP -- I'm just going to use ISPs as an example, when your ISPs receives that subpoena, they generally notify you. You usually get about ten days to two weeks to try to quash that subpoena, if you choose to do so, and if you can get a lawyer.

One of the problems that we found again and again is that generally these cases do not get brought in the area where the person is whose information is being subpoenaed, so like say I live in Pennsylvania. I

posted something on a Yahoo! message board, so the subpoena is served in California where Yahoo! is, so it's very difficult for me living in Pennsylvania to secure legal representation in California to try to quash that subpoena.

So it puts a great burden on the person who is trying to engage in speech, and indeed I think does in many ways get people into not speaking because they fear this kind of legal process.

MS. ROBBINS: Before I get to the question in the back, there's a question card that I have that addresses that issue that you just talked about, Annalee.

The question card says: "It is important to note that authentication of an email address does not necessarily imply the authentication of the individual using that address. I can get an anonymously named account at Yahoo!, but the email account will be authenticated coming through Yahoo!"

Is that true?

MS. NEWITZ: You can't get an anonymous account with Yahoo! This is a question that comes up a lot. People will say to me, "Well, why don't you just -- if you want to speak anonymously, why don't you just sign up at Yahoo! or Hotmail with a fake name?," so the

question is, do we really want to make honest people dishonest in order to speak anonymously, and I say no.

MS. ROBBINS: Do you want to clarify?

MR. ANDERSON: Dave Anderson, A-N-D-E-R-S-O-N. The forensics that are available using IP addresses today, Annalee, are such that you would have to have a real incompetent attorney to not be able to figure out who you were based on spoofing. If there are not other mechanisms created such as sites or such as ISPs that will not allow you to track back, you're going to get found out very easily, so I would suggest authentication isn't going to change that picture much at all.

MS. ROBBINS: There's a question back there on the left.

MS. GRANT: Hi, I'm Susan Grant from the National Consumers League. We've heard about the intangible costs of authentication in terms of the potential to chill free speech and discourage whistleblowing. Can any of the panelists comment on potential tangible costs to the end user, either directly or indirectly, for the ability to authenticate or for the ability to remain anonymous and what impact that might have on individual users, small businesses and small organizations?

MS. ROBBINS: Jonathan, would you like to

answer?

MR. ZUCK: Sure, I'm happy to address that. I think the tangible costs to consumers and small businesses would be a negative one. I mean, the bottom line is that the costs associated with spam and with online fraud in the form of phishing and other vehicles is so high right now that everyone is clamoring for some kind of solution. There's not an implementation clause for a particular end user or a small business to have authentication in place.

This community instead is spending millions and millions of dollars on their own little versions of filtering software or whitelisting or blacklisting and trying everything they can to spend whatever money they have to try to stem this problem.

So the bottom line now is that while we've had this panel, 200 more messages have arrived in my inbox telling me things I need and somehow both Citibank and EBay have lost my password in that time frame as well.

So the bottom line is that the real costs are associated with the problems being addressed, and the costs that will be born through an authentication system are going to be born by the huge ISPs and others that are going to be doing that authentication on behalf of users, and they're already bearing huge costs in the

form of filtering out as well.

So everybody will save money and increased productivity I think with authentication in place.

MS. ROBBINS: I think Duane wants to also respond.

MR. BERLIN: One example of a cost that's currently being borne is the lack of an effective way to deal with authenticated emails is a number of legitimate senders of commercial emails that do not hide their identify, do not engage in any other practices that are within the commonplace menu of the spammers are being blocked by the ISPs for various reasons based on voting or imprecise internal standards that the ISPs themselves implement.

And these are a tremendous cost to the small and mid size businesses that attempt to use email legitimately and aren't trying to hide their identities so a reconciliation of the process that is aimed at those that are specifically trying to hide their identity would bring tremendous savings to those businesses who are trying to engage in legitimate commercial speech and really on a practical level being deprived privately of their ability to do that.

MS. ROBBINS: There's a question all the way in the back by the door.

MR. BAKER: Phillip Baker with VeriSign. Thank you very much for holding this meeting. Point to Dan. I was with the web team when we were having the fight with Gofer. The thing that actually killed Gofer was when the university for which Gofer originated decided to start exercising copyright over the Gofer code, and that was what killed them. That allowed us to beat them, so you actually were making a worse point than you could have there.

The point of the GPL was it came out of an era where university copyrights would be public, with public money and then turned into private property somewhere along the line in a very suspicious way.

I think what we've got here with the patent issue is very different. Patents are a very different form of property and trying to squeeze everything into the GPL ain't going to work, but the other thing that doesn't seem to be working is the sub-licensing issue, and in particular this whole myriad of bilateral agreements that you seem to be getting worried about, that if I have to have a bilateral agreement with Microsoft and Intel and everyone of the other 50 potential IP holders that might be involved in a moderately seized IP.

So maybe what we need to do here is to change

the model, and there is actually a legal model in existence that's being used in other forms, and that's the rule book. If you joined the stock exchange or metals exchange, you sign the rule book of the exchange, and then you agree to undertake certain -- you make certain undertakings that you agree to and you get certain rights back in return so instead of having to call a contract with every other member of the exchange to say you're going to recognize the contracts, you have a contract with the exchange, and then the exchange has the contract with the other person.

Maybe now that the MARID thing is kind of settled down and people have stopped getting quite so paranoid about the situation, maybe we could look at that type of model and maybe get that situation, because Scott is right, what's happened here is not that Microsoft did something that was unusual or something that was completely out of the ordinary.

What happened was that Microsoft did something, and people noticed that, oh, look, the whole open source movement and the whole movement and the whole way that IP is being licensed, those two do not add up. There's inconsistency, and we need to get over it, and we can't get over it by simply stamping our feet and saying, it has to be my way or the highway because you don't have

the patents.

MR. QUINLAN: I think the analogy made as to why GPL is a good one because we have a similar situation with Sender ID where the SPF standard was out in the open by the open source community, and in essence a company tried to take it private by adding a portion of their own technology to it that wasn't encumbered beyond what the original specification was, and that's why SPF is free to use for everybody and PRA is not.

MS. ROBBINS: We have time for one more question, the gentlemen with the beard.

MR. HAMMER: Michael Hammer, H-A-M-M-E-R. I did participate in MARID and the SPF group and what not. First off I would like to say this is really about open standards, not necessarily open source, and one of the concerns that I had when MARID was dissolved, the indication of my ATF was go out, submit the drafts as experimental, let's see what works out in awhile.

Now, SPF was against public records on SPF 1, and when people put those records out there, what they were really doing was making a claim as far as the RFC 2822 mail fraud, the domain.

Recently Microsoft has unilaterally decided not to apply PRA against SPF 2.0. Instead they're claiming it against SPF 1 records. This breaks the intent of the

publisher of the records. It causes legitimate mail to be rejected, so my question would be for Mr. Kaefer.

Why did Microsoft decide to apply these checks against SPF 1 knowing that it would break the intent of the publishers?

MR. KAEFER: I have to admit this is one of those cases where I'm not an expert, but we have one in the audience, and if it would be okay, we'll have Harry respond to this.

MR. CASE: My name is Harry Case, and I work on the technical aspects of Sender ID for Microsoft, and I wanted to address the issue that has just been raised.

First of all I want to point out that we did not unilaterally decide to make this decision. There was some significant discussion about this in the MARID working group and indeed afterwards, and the very strong feedback we got was that it was important to preserve backwards compatibility with domains that had already published SPF records. That's the first point I would like to make.

The second point is that we've looked at this fairly closely, and we believe for the vast majority of domains that published SPF records, that the content of that record would be identical regardless of whether the Mail From check or the PRA check are being implemented,

and rather than impose the requirement on all domains to publish two identical records in the DNS, we felt it made far more sense and was far more efficient to simply have one record that is used for both checks and provided provisions or mechanism for domains that do need to make distinct records for each check available, so they can do that if they need to but that's on an exceptional basis.

MS. ROBBINS: I want to thank all the panelists for joining us this morning. I think that we've had a really rich discussion about these issues, and we are again now going to take a break. There are refreshments in the back. I believe there's coffee and bagels and all sorts of goodies.

Thank you, and we'll see you at 11.

(Applause.)

(Break in the proceedings.)

PANEL 2: EMAIL AUTHENTICATION PROPOSALS:
CRYPTOGRAPHIC APPROACHES
MODERATOR: DONNA F. DODSON, NIST
PANEL MEMBERS:
MILES LIBBY, Yahoo!
JIM FENTON, Cisco Systems, Inc.
DAVE CROCKER, Brandenburg InternetWorking

MS. DODSON: Good morning. My name is Donna Dodson. I'm with the National Institute of Standards and Technology, and we, at NIST, are very pleased to be co-hosting the E Authentication Summit with FTC today. It's delightful to see so many people participating in this, and I think the morning session, the first session, really set up the business requirements and some of the privacy issues and some of the legal issues that we need to think about as we move forward with dealing with the problem of spam and email.

What we're going to do in this particular session is to look at three technical proposals and have an understanding of some of the technical options that are out there. In particular these three technical proposals deal in some very different ways, but have an underpinning of cryptography with them, and as everybody knows, we used to think of cryptography as being

primarily used for confidentiality, and in today's world of business, the use of cryptography for integrity and authentication really has taken precedence and has taken hold.

There are always some understanding when one starts to look at cryptography with scaleability, intra operability, all those good words when you look at something as massive as email.

We're going to run this session a little bit differently than the previous one in that we have three proposals today that are going to be discussed. These are all three RFCs that will be discussed, and I would like to hold audience questions for the last 20 minutes of these presentations, so that you can get a feel for the technical underpinnings of each one of them.

I think it ought to create an opportunity to ask some very good questions and compare some of the different solutions that are presented today, and I would like to ask each one of the speakers at this session and the following sessions to really make sure that they're using the microphones because people in the back have said they're not able to hear.

If you're not able to hear us, please raise your hand, and we'll use that as a signal, not as a question, but as a signal that you can't hear, so you're able to

pick up everything okay? Very good.

All right. Our first presentation today will be on DomainKeys by Miles Libbey from Yahoo! Mail, and with that, I'll let you get started.

MR. LIBBEY: Good morning. I'm Miles Libbey. I'm the Anti-Spam Product Manager for Yahoo! Mail, and I am going to talk about DomainKeys.

When we started thinking about sender authentication, we reflected on our experience in Yahoo! Mail. We've been running a reputation engine in Yahoo! Mail as part of our anti spam efforts for the last five years, launched in 1999, and it's based on IP addresses, and we found that IP addresses are really insufficient for email identity. They don't work well in a number of cases.

First, they don't work very well with the email service providers. This is a case where a company outsources their email sending to aid another company that specializes in email sending. So when a company does this, and ESP sends mail to these other companies, they frequently consolidate all of their sendings through a certain small set of IP addresses, and this makes it hard for a reputation engine to determine the difference between the reputation of one sender versus another.

Similarly, IP addresses don't survive forwarding, so when EBay, for instance, sends a mail to somebody who forwards their mail, when the end recipients receives the mail, their reputation engine thinks of the mail as coming from the forwarding mail system, not the initial author of the mail, and since forwarding systems generally forward all mail, they end up having a very mixed reputation.

Some of the mail will have very good reputations and some will have very bad reputations, but by using IPs, the reputation systems aren't able to distinguish between the two.

Finally the IP addresses are invisible to the user for the most part. They don't know or care about IP addresses, so when we think about reputation systems, we think about using the domain, typically the frontal domain in the body of an email.

So the DomainKeys technology is actually pretty simple. First what happens is the domain owner self generates a public and private key pair. They then publish that public portion of that key to a new standardized DNS text record. The public private keys are solely determined by that domain owner, and this DomainKeys is actually just as secure as DNS, so many, many users and companies are using things like Web

Services Today. DomainKeys is as secure as that.

The DomainKeys then -- domain owner then can revoke the domain key as well, and actually the DomainKeys allows for the domain to have multiple keys per domain, so this enables a domain to give out a key to an ESP, so you can have multiple identities. You actually can trace a particular key to a particular user name, and if you were to give out a key to an ESP, you can only revoke that key after your contract is finished.

So once you've generated then the set up portion, then it's time to move on to something you can verify, so outbound email is signed with this private key, so you put the private key into your mail server software. The mail server software performs a mathematical algorithm and generates a digital signature which then is put into the header in the email.

The digital signature covered the headers of the email as well as the body so the actual DomainKey header actually adds about 150 bytes to a message.

Then the email send off is normal, so when the receiving system receives that email, they find the domain in the body of that email, lookup the DomainKey record from the DNS record and verify -- perform another mathematical algorithm with the signature, the DomainKey

and the content of the email. They can run this mathematical algorithm to verify this message was indeed send by the author of the message.

So we really designed DomainKeys with flexibility in mind and trying to minimize adoption hurdles, so we wanted to reuse existing hardware such as DNS and software to really minimize the deployment costs, so we're using standard sign-in technologies such as RSA. This also enables us to use other technologies or allows other technologies to use digital signatures and cryptography in the future.

For instance, in the future one might use the Bounce Address Tag Validation with DomainKeys. By using DNS to distribute the keys, I think that this enables the use of DNS caching to have significant performance benefits.

So there's a number of use cases that come to mind when examining any center authentication technology. One is when an ESP sends mail on behalf of a company, so the company could give their ESP a private key to use for sign-in, and obviously publish the public portion of that key into their DNS record. You can actually train that key to be used for a particular user name, such as sales ad or marketing ads, and then once your contract is done, you can revoke it or you can

revoke it for any other reason.

You could also delegate your subdomain of your DNS record to that email service provider, and this will give the service provider responsibility for managing the DNS as well as the mail server software, and again you can revoke that delegation at any time.

Another use case is the mailing list for discussions, so there are generally two cases in mailing lists. One is that for mailing lists that don't change content, so in this case the signature is generally not broken, and you can -- the receiving system can verify that the original author sent that message, so the mailing list can actually choose which reputation it wants to apply to its email, whether it wants the reputation of itself applied to the email or the if it wants the reputation of the original author of the message applied to its email, and it can choose whether to sign the message or not.

For mailing changes that do change the content such as Yahoo! Groups, which has an ad in an -- advertising in an unsubscribed options at the bottom of the email in terms of subject there are a couple options for them. One option is to add a sender header, and thus take responsibility for that message, and then resign the message and claim it from Yahoo! Groups, for

instance. This actually is likely what the ISP wants the group to do. They want to be able to apply the reputation of the mailing list to that email.

So another case in the email world is in forwarding. Forwarding is actually quite simple in the DomainKeys. The original author signs the mail using DomainKeys and the message is verified using DomainKeys.

Another use case is when various web pages have news pages such as send this page to a friend, so if you're on the New York Times web site, for instance, you can send this message or send the page as an email to somebody, so the news source can also claim authorship of this mail. They have a number of options as well. They can set the from address to be "news articles at New York Times," for instance, or they can set the adjustment to be the user on the computer and put the sender address to be "news articles at New York Times" and claim authorship or sign the message and claim authorship that way.

They could also set the reply to header as the user's address so that if the recipient actually clicks on that reply then the message will end up back at the user that initiated the sending.

So we could also talk about licensing for DomainKeys. Yahoo! has filed two defensive patents

surrounding DomainKeys. Our patent license is really designed to allow freedom to operate while protecting the industry, so the highlights of our license are that it's royalty free, it's sub-licensable, and it's perpetual unless you sue Yahoo! or other implementer over DomainKeys.

We also do not require any registration. You're granted a license by simply copying and pasting the license there.

So there's a number of issues that come up with when you're talking about any cryptography solution. One is CPU cost. Sendmail actually did a study on the CPU cost this past June on DomainKeys, and they found that running DomainKeys added between 8 and 16 percent CPU cost to mail server software. So this was actually quite nice.

Generally mail server software is not CPU bound, and we feel like this kind of CPU cost is not going to be significant for the vast majority of senders. What is CPU cost? It would be by doing a cryptographic signature, the mail server software -- hardware needs to perform additional operations, additional mathematical operations, and this is CPU intensive or could be CPU intensive.

Another issue that comes across frequently is

one that is a replay, so this is the case -- so while DomainKeys enables forwarding to exist spammers could potentially use this against us, so a spammer could sign up for a free service such as Yahoo!, send themselves some mail and replay that message off to -- and send it over and over and over again to lots of different people.

This is not really an authentication issue. It's more a reputation issue. Once Yahoo! has enabled a user to Sendmail. We are in fact claiming the mail is coming from Yahoo!, so by replaying your own identity, you can ruin or harm the reputation that you already have, but the original message was authorized and you can't change it in any way, and you can't change -- you can't replay a message from high value identity mail such as EBay or Citibank or what have you.

Another issue is that of message integrity. So when the message is signed with DomainKeys, we are protecting both the content of the email, we were saying this email is indeed created by the author of the message as well as it came from this person.

So small changes to the message will invalidate the signature, and say if you add text to the bottom of the body, no longer will the message be authored by the original sender. You need to -- the DomainKeys check

will begin.

So one solution to this is that whenever changes to the messages are being made is the changer can actually resign the received message and thus claim ownership of the mail.

So DomainKeys, it was submitted to IETF. The latest implementation was sent to the IETF in mid August. Yahoo! Mail is in the final stages of deployment today and SBC, British Telecom, and Rogers implementations will follow shortly. Similarly, for verification, Yahoo! Mail, SBC, British Telecomm, Rogers will all begin verification deployment very shortly.

We're also receiving reasonably strong industry adoption. GMail has already begun signing all its mail. Sify last week began signing its mail. ISP in India, SkyList. A direct mail ESP has begin signing, and AOL and Earthlink have also indicated their interest in testing.

We have released a royalty free open source reference implementation of DomainKeys on source forge to enable other MTA developers to have an easier job of implementing DomainKeys.

Today, Sendmail, Key Mail are proposed actively using DomainKeys. There is an exchange version that's coming out from CERN, the specific one that created

the Internet. Several other commercial or mail server software systems have announced support such as Port25, Omni IT, E-Type and Active Software.

So you can find more information about the specifications on the Source Forge site DomainKeys.SourceForge.Net.

Thank you.

(Applause.)

MS. DODSON: Our second panelist will be Jim Fenton of Cisco Systems, and he's going to be talking about an RFC Identified Internet Mail or IIM. I keep writing it down IMM. Sorry about that.

MR. FENTON: Good morning. I would like to talk to you a little bit about Cisco's message signing proposal Identified Internet Mail, and I'm going to talk to you about it mostly from the standpoint of what it means to users of email and to administrators of email domains that would be involved in using it.

Let me start by talking about sort of what we were trying to accomplish with Identified Internet mail. We began with the notion that we shouldn't break email as a whole. The reason that we have the problems that we have is because email is a very successful medium. The spammers wouldn't be using it if that weren't the case.

So we want to keep the positive aspects that people can send to others without introduction. There is some degree of anonymity. Of course there was a lot of discussion about that earlier, but to the extent that there is anonymity now, we wanted to preserve that, and we want people to be able to continue to send email independent of where, in the Internet, they're submitting their Mail From, so if they've traveling, perhaps they want to be able to send messages from their hotel or from their cell phone or their PDA or something of that sort.

So our goals are here as much social as they are technical. When we talk about anonymity, we believe that there's a strong distinction between a message being anonymous and a message being fraudulent. If somebody wants to use an unknown email address to send a message to someone else, I don't have an objection to that. I do have an objection if they use my email address to send a message to someone else.

We want to try and help this get adopted by requiring as few changes in the infrastructure as possible, so some of the things that might have to change are mail servers known as mail transfer agents and operators of mailing lists.

We want to try and make it so that if a message

ought to succeed, we want to try and find a way for messages to pass authentication than to try and find a way to throw them away, so we're trying to accommodate the common behaviors of mailing lists and mail servers, and we're also trying to help adoption by making the trust base for this something that's very light weight, already existing to a large extent, ultimately based on the Internet domain name system, but yet something that will scale very well, and we'll talk about the importance of that later.

Finally we think that the ultimate solution to these problems is going to involve accreditation and reputation systems, and we want to have something that is reliable enough as an email identifier that people can have something to base those systems on.

So our model is maybe a little bit different than the authentication as it was introduced by John Levine earlier. We consider that there are two parts to what we do. We authenticate the message, and note that I'm saying authenticate the message and not authenticate the sender.

We're trying to determine that the message that we received hasn't been modified in transit, maybe subject to a cut and paste attack or altered by -- well, just some ownership server transmission problem, and

then the second part of this is that we want to determine whoever it was that sent it, we're not asking who it is, but whoever it was that sent the message we want to determine if they were authorized by the people that ran the domain.

We consider the addresses to be the property, if you will, of whoever is registered for that domain, so the administrator of the domain should have the right to delegate that authority to individual users.

People have a tendency to confuse email addresses with identity. They're not the same thing. People do change ISPs. Addresses get reassigned to different people I'm sure. People change companies, and just because you have a particular email address at a particular time doesn't mean that you will always have that address or that authorization from that domain, and it also doesn't mean that the domain administrator, if they really wanted to, couldn't appropriate that for some other use.

So this is a diagram of sort of a typical message flow. There are lots of variations on it. Signing and verification messages can happen in different places in the system, but I'm kind of illustrating the common case here. Someone submits a message through their own originating domain. A mail

servers does the signing. They don't need any new software on their PC or whatever.

It passes through the Internet to the recipient's domain. A mail server does the verification there and consults with the originating domain to find out whether the key that was used to sign the message, which is sent in the message in our case, whether that key is authorized by the originating domain to be used with that email address, and if both those tests pass, then normally the message is marked to indicate they passed the test and passed the recipient.

In the longer term, the recipient domain can also apply some of their own policy. One of the important aspects of our proposal is that there's the ability of a sending domain to publish a policy that says, we sign a hundred percent of our mail messages.

If you receive an unsigned message that is supposedly from us, it's probably not something that you should trust, so it supports the anonymity by a domain that doesn't have that policy. People can send messages unsigned, and they'll be treated in some manner by the recipient, perhaps not sorted into as high a priority mailbox as signed messages, but when there's a policy from the originating domain that says, we intend to sign all of our messages and the recipient gets one that

isn't signed, they can do something that's somewhat more drastic in terms of either warning the user or potentially even blocking the message in order to respond to that policy.

Finally, we expect that reputation and accreditation services will be part of the framework and although we aren't addressing it in our proposal directly, we think in a very similar sort of transaction as you do with the originating domain could be done with the reputation service in order to find out something about the originator's -- how you should deal with that particular address that you've just verified.

We think that there are a lot of cases where user level keys will be important, and I want to distinguish user level keys as distinct from user identities like I mentioned earlier.

We think that there are a lot of use cases where people will need to sign their own messages. People will send their messages from outside their domain. They'll be traveling. There are some domains, a lot of you probably have college alumni associations that give out email addresses, and those domains might in the long-term want to support the ability for you to send your mail directly as you do now, maybe with some software on your PC that does sign-in for you, and don't

have to route it through the college or organization of whatever sort.

When you have these sorts of capabilities, you want to operate on the principle of least privileged. You don't want to give people authority, a key authorization if you will, that will allow them to do more than they ought to do. I wouldn't like everyone that went to my college to be able to send email as any address at the college.

Likewise, if I was a company that wanted to contract with a marketing partner to conduct some sort of an email campaign or perhaps to send benefits messages to my employees, I wouldn't like to -- it requires a higher level of trust if I was to give them a key that was authorized or for them to generate a key that I authorized that's authorized for any address in the domain.

It helps the relationship, it requires a lower level of trust if you can give them a key that's more specifically authorized.

There are other situations like that where people need to have the ability to send email on behalf of others. An administrative assistant might have several people that they send email for, on behalf of, and that assistant would like to have the ability to use

the same key all the time and just have that authorized for multiple email addresses.

There will be -- so we expect that a few domains or quite a few domains will need some user level keys. A few, but some, will need large numbers of keys, and we have to provide the key authorization for those domains to scale to large numbers.

So here's a little more discussion about the use cases that we're considering. We're approaching this problem both from the standpoint of our customers that are enterprises as well as our customers that are services providers.

I mentioned a minute ago that you can contract with a third-party company to authorize sign-in. There are quite a few cases where employees that are distributed around the world as they are these days want to be able to send business related email directly from their home or remote office or whatever without requiring to do a VPN connection with the parent company.

Mobility is also getting to be one of the things that's scaling up, and to the extent that you need user granularity keys to support mobility, I think that need is going to grow and is something we need to plan for in our scaling.

Mailing lists can do a lot of things to messages. We're trying to handle the common cases like changes to the headers and messages that are appended to the bottom and allow those messages to flow through unmodified mailing lists. In the longer term, we really expect that mailing lists will sign messages on their own behalf, but in the meanwhile we would like to have mailing lists work on a best effort basis.

I mentioned affinity email addresses so these are like college alumni associations, organizations like IEEE, other professional groups, hobby groups and so forth. Users will have multiple devices that they send messages from, so sometimes they'll use their PC, sometimes their cell phones, sometimes their PDA, and we need to have the kind of scheme that supports that as well.

And I think Miles mentioned mailing a news story to a friend sort of thing, the third-party message transmission, which is a common case. Another is invitations, EVites, things of that sort, where the service depends on the ability to send mail as the customer, if you will.

So here's my one geek slide I guess. This is an example of what the message headers for one of our signed messages looks like. The content that's in

yellow there are the elements of the signature. We include the public key in the message because it's an easy way of distributing the key, and it allows us to do some checks even without checking with the originating domain.

The signature is computed over the content in the message as well as selected headers that are specified by the originator, and then finally we have copies of the headers that we're signing, and we include those in order to improve the resiliency of Identified Internet Mail against modifications that mailing lists and things of that sort might do.

So that the message even if the -- for example, the subject of this message had been modified. The recipient would be able to replace the original subject or just flag that the subject had been modified and still accept the message, so that's one of the efforts that we're trying to make in order to improve the verifiability of messages that go through this.

So a lot of things have changed since Internet mail was defined. John Levine talked about the difficultly of layering trust on top of something that was designed without it, and we think that what we've done here is a good trade-off between being a complete solution to the problem and something that's exceedingly

complex. We're open to working with others in order to further refine this.

Thank you very much.

(Applause.)

MS. DODSON: In our third presentation today, Bounce Address Tag Validation will be given by Dave Crocker, Principal of Brandenburg InternetWorking Group, and I just think it's very interesting the differences in approaches that people have taken and some of the similarities, and I think we're going to see that a little bit more even in the third briefing.

MR. CROCKER: Thank you, Donna. Good morning. It's a pleasure to be here in spite of the motivating cause. The FTC Workshop that was held about a year and a half ago on spam seems to me to have been a seminal event in terms of discussion on this topic. I'm hoping that this event serves the same purpose with respect to one aspect of pursuing that, and what I'm going to talk about is a proposal that's independent of the two that you've just heard, although it can serve as an adjunct to them. It uses encryption to do signing as they do, but in a very different place.

With respect to most spam control techniques and especially any that purports to do authentication, what we're finding is the first and I think most important

step is to decide precisely what you're trying to achieve. Signing can be done in many places, in many ways, by many agents, and so we need to be very precise so that there's no confusion about who is doing the signing and what it means to do the signing.

That's what the subtitle on this is trying to answer with respect to BATV. I should comment that BATV is a collaborative effort. There is a design team that works on both BATV, and you'll hear about CSV in the next session, and in fact, it comprises the authors of those two papers, those two proposals and a couple more people. The design team is mostly occupying the front row in front us today here, so there will be an easy ability to clarify any confusion that I create.

There we go. So by way of showing that there are many possible agents that can do signing or otherwise take responsibility, in a typical email, and this is not a complete list, it's just a useful subset, there are five different entities to be aware of in terms of basic roles, and the distinction between the originator and the submitter or what in RFC 2822 parlance is called the sender, is an important one.

One that is responsible for injecting the message into the service and the other is responsible for creating the content. The BATV focuses on a

different string, and the best term for that string I think we're finding is to call it the bounce address, but unfortunately what it's called in RFC 2821 or SMTP parlance is Mail From. We goofed. We didn't really understand what that string meant, and what is amazing is it took us 25 years to find out that we made the error.

The string does not have to bear any direct relationship with the from or the sender field, and in fact in many very legitimate bulk sending situations, it is completely independent because you want to direct bounces to a special bounce handling facility.

So the purpose of BATV is to sign the Mail From field. Why, why care about signing that field? And the answer is something that everyone in the audience already knows, but I'm obligated to go through a list.

One is that spammers by way of making their life more convenient direct bounces somewhere else, away from their sending environment, and as I just said, legitimate bulk mailers often do direct bounces to a different location, because for any large email sending situation, the number of bounces often is quite large, and the handling of that stream can be its own speciality, and so spammers just move the problem to someone else, like you or me.

The other is that this has become a very effective technique, the sending of bounces or messages appearing to be bounces as a back-door Trojan into your machine where you handle it differently than you might handle a regular piece of mail, and then lastly, because of that first bullet, that's a flood of messages, and that's called a denial of service attack hurting your capacity.

So just to make sure we understand the sequence of handling in emails, somebody sends a message, and it gets to an MTA which tries to deliver it. A mail transport agent tries to deliver it to a delivery agent, and the delivery agent says, "No, you can't do that, I don't have that address," at which point the MTA then wants to generate a bounce, and they send the bounce back to the bounce delivery agent, so that the entity that creates the bounce message and the entity that tries to deliver the bounce message are the two most interesting in this scenario.

What BATV does is with respect to that last step, the bounce delivery agent, the question is, should I actually deliver this to the user because if this isn't really a valid bounce, it would be helpful for me to not burden the recipient with this traffic, and all of us I think get highly distracted by the receipt of

all of these invalid bounces, and so it would be nice to have that filter.

It doesn't save any email infrastructure resources, but it saves the recipient of the bounce, and that's a nice thing to do. Even better would be if the entity that's creating the bounce could decide not to do that, if they had some way of going -- some way of saying, I believe that this bounce address is invalid and therefore I will not send a bounce, and that will save an enormous amount of Internet mail resources.

It turns out that capability leads to an interesting additional one, which is if I know that this is an invalid bounce address and I can determine that early in the transmission sequence, I probably have a message that isn't valid so I can use that to decide not to send the message itself further on, and that would save even more resources.

So how does BATV go about doing this? It puts a signature onto the Mail From field. BATV is in fact a framework for different signing techniques, and the reason that we took the approach of having a framework rather than a single technique is the world of signing and sealing technologies seems to change quite often, and there are different approaches. There are different needs for different users.

So we decided not to try to claim that we knew which ones would be a winner, but rather to allow a -- to allow a standard framework that would allow encoding alternate approaches and let actual usage determine which ones became popular.

The approach to encoding the signature is to take the existing Mail From field and add two fields to it, the one on the left called sig-scheme is just the name of the approach to doing the signature, and then after the slash, it's follow the data, whatever is appropriate for that signature technique.

It turns out that while most of what we think about for doing signing in the public world is using public key technology, this is one very unusual environment where private key technology actually is viable, and in fact in some cases it's better, and you get different benefits from using either of them. The main benefit of not using public key is you don't have to worry about a public key infrastructure.

It happens that since the recipient of the bounce is in the same administrative domain as the sender of the bounce address, the originator of the bounce address, there is some chance that you can administer private keys in a way that works.

The specification for BATV includes one

technique. It's the simplest one we could come up with, because it's the one that John Levine is already using. John is one of the authors of the BATV, and in fact this is all based on his idea.

Signing the Mail From field or authenticating the Mail From field is something that people have been wanting to do for awhile, and this technique doesn't require registering a path all along the way, so when we mentioned it in a conference presentation, it struck me as just the thing that we ought to try to pursue.

The public key approach uses the same basic style as the private key approach, but the difference then is that we need a public key infrastructure. You heard some references to that in the previous two presentations, and the assessment the design team made with some pain, because we really like the idea of a public key approach, was that the critical part of doing a public key approach is the public key infrastructure.

Creating a public key infrastructure is difficult and it would be foolhardy of us to try to propose a solution to that given that there is a long history of trying to and not much success so far and what we want to do is tag along -- notice that address tag along validation. We want to tag along with any existing public key work that gains any popularity.

So an example of that would be public key mechanisms that are based on the DNS that you've heard proposed in the previous two presentations, and if it turns out what that leads to if you use an IIM or DomainKeys is that the signing of the Mail From let's you do an envelope time or a reception time, preliminary evaluation of the overall integrity or validity of the message where you can save the deeper analysis for the time you're looking at the internal content.

Because BATV focuses on the Mail From, it's worth paying some attention to alternate techniques for validating the Mail From, and I characterize the approaches as one being object based which is BATV and the other being channel based, which requires that you register the transmission path, so the object approach for BATV says we're going to wrap up the sensitive data, and then we don't really care very much what path it goes through, if it goes through a path.

We wrap it up, and then we go through whatever path we want, and this slide will show the recipient, but it could be an MTA along the way that does the unwrapping. We're insensitive to the intermediate points.

Path registration doesn't wrap up the data, but rather tries to protect the entire channel, and it does

that by having the originator register the paths that the message is going to go down through, and if you have a path that isn't registered, it means that the recipients down that path don't get a protected message. They can't certify it, and you have to go back and fix that before you can certify those additional recipients.

Status of the project? Let's turn to that there. We've gone through a couple of rounds of specification, a whole lot of public discussion. I would say that the specification for BATV is in a pretty good state. To my knowledge we only have one deployment which is John Levine's, and he hasn't upgraded the syntax yet, has he?

No, not yet, so he's been using his original syntax, and that's an important difference for the public interpretation of the format, but it's not important for the semantics of the proposal.

We're looking for people to test this. The neat thing about testing the private key is the only people who have to adopt for you to get your benefit is you. You don't have to have me or any of the rest of us adopt your change. As long as your originating site that creates the bounce address and the sites that are referred to by that bounce address collaborate with each

other and they presumably are under identical administrative control, then you will get the benefit that you are looking for.

We are in the process of pursuing IETF working group status, and that will proceed in the usual fashion. We have a draft charter, and we have a discussion mailing list that covers both this BATV and the CSV proposal you're going to hear about.

Places to go for the mailing list is at the MIT Association site, and these specify the proposal itself is the mass BATV. There's a larger framework document that tries to provide some standard terms of reference for email architecture, which is also an Internet draft.

So none of these documents have changed the stable publication of RFC, Requests For Comments, which isn't the Request For Comment, but they're in the Internet draft stage, which is the request for comment.

Thank you.

(Applause.)

MS. DODSON: We have an opportunity for questions, and we have some microphones available, so if you want to take the microphones, and if people could state their names, spelling of the last name, and then your question, please.

MS. BAKER: Hi, my name is Dawn Rivers Baker.

I don't really have to spell that, do I? This all sounds very tidy in terms of the way you're envisioning people using email. I'm thinking of a scenario where if I want to send email from my domain at MicroenterpriseJournal.com, that's fine, I have the domain name, and I send it through my pop account, but if I want to send an email from Dawn at DawnRiversBaker.com, well that domain is parked somewhere, and when I get email to that address, it's forwarded to me, and when I send email from that address, it doesn't go through DawnRiversBaker.com.

It goes through my ISP at my house, which is RoadRunner, and would this system accommodate all of this?

MS. DODSON: Can you hear me? Which system are you looking for.

MS. BAKER: In other words, would the cryptographic systems at any or all of them that we've just heard discussed be able to accommodate somebody using email without using a pop account where they use email forwarding to and fro and where they send out through their home ISP as opposed to a pop account?

MR. FENTON: Sure. Is this working? That's one of the benefits of the cryptographic system is that you -- it sounds like you want to be able to send mail

from an arbitrary place. It may always be your home. It may not, or in some cases your home ISP or your address on that network may change from day to day, but in this case it would require some software on your PC because you want to sign your mail directly.

And we expect that software to be developed, but that's the beauty of this is that really it sort of follows the postal model of drop the letter into any mailbox in a sense.

MS. BAKER: Thank you.

MR. LIBBEY: I would also say it's possible that your ISP could sign mail for you. You could give -- as the administrator of your domain, you could give your ISP a key for your domain and have it sign for you.

MR. CROCKER: I think there's some potential confusion because both of the other proposals focus on what will be the original implementations which is through the MTA. My experience says that when you do an architecture that requires the use of the infrastructure within the scheme where MTAs are part of the infrastructure, when you do an architecture that requires that, there's massive burdens for large scale adoption.

That's different from having an architecture which is really defined in terms of the end system and

can be implemented in the infrastructure for convenience, and that's what is true in both of these proposals.

In point of fact you can have user agent software implemented and the MTAs don't have to know anything at all about it. However, it's convenient especially for large ISPs or any other enterprise service situation to have the MTA domain.

MR. LEVINSON: Andrew Levinson, L-E-V-I-N-S-O-N. The public key proposals have both CPU costs, which Mr. Libbey mentioned but also have costs in the use of the DNA. Do you have any estimates on the load on the DNA system? I'm sorry, DNA -- DNS system. Thank you. I guess I'm a little nervous.

So the cost in the DNS system for sort of public key implementations?

MR. LIBBEY: So certainly for every single email sent today a DNS lookup is performed to find the MX record, and all these DNS lookups are indeed cached by the vast majority of implementations, and this would be very similar in the case of I think all of these proposals, so the recipient system would do a DNS lookup. It would cache that result until the next time you send the mail that would not require another DNS lookup.

Today's mailing systems frequently do many -- other DNS lookups such as reverse lookup, such as MS lookup or call backs, what have you, so we don't think this is a major burden for MTAs.

MR. FENTON: There are actually two sorts of costs. One is the number of lookups that you do, and the other is the size of the lookup. Both of the proposals support doing -- basing the trust on DNS. We use it in different ways. DomainKeys retrieves the keys from DNS, Identify Internet Mail, it just checks the authorization of the key by DNS, which is a somewhat shorter transaction, but both of those can be cached.

Where the caching doesn't work as well is when you have large numbers of individual keys, and in those cases, Identified Internet Mail has a second method that can be used, which is to use -- it's actually a web server sort of based piece of infrastructure that we created called a key registration server, where all the DNS would have to do is find the location of that, and then you do a separate transaction, which can be cached directly by the verifier in order to determine the authorization of the key.

MR. CROCKER: I'm really glad Ed asked this question because the query cost when you're crossing the Internet half way across the world is a non trivial

point to pay attention to, and there are some proposals I think which have some unbounded costs there, but the encryption based ones I know about all have pretty modest costs, essentially at the level of one or at most two lookups.

So while there are some people who are concerned that the aggregate use of these techniques on the whole Internet will have a problem, it doesn't really look like that's a concern.

MR. GILLUM: My name is Elliot Gillum, G-I-L-L-U-M. I think there's some important aspects to the CPU network costs in the various proposals that are easily overlooked. For instance, when you're signing your mail, the outbound mail systems are extremely unlikely to be CPU bound so the CPU cost is likely to be negligible, and on the inbound systems, to the extent that you can mitigate the cost of what's likely your only CPU cost which is spam filters, you can actually make up possibly or gain whatever cost you're paying in validating signatures.

On the DNS stuff, there's actually -- you still have to transfer the key and the authorization, so you may be paying the cost in transferring the email versus in the DNS request, so you do gain some benefit in terms of the size of your cache, but you are still paying for

costs in transferring the key and the message or in the DNS, and there's a subsequent cost in storing that key in the message in that proposal.

MR. CROCKER: This was labeled a technical conference, wasn't it?

MR. QUINLAN: Hi. Daniel Quinlan, Q-U-I-N-L-A-N. So my question is more so directed at BATV because the other two proposals don't have this issue, in that when you send a message, you decide to sign a message with IIM or DomainKeys, then there's no real effect on whether your message is going to get delivered or not whereas with BATV, there's at least one case, the curiously named easy M-O-M mailing list software where it would use your Mail From address, the bounce address, to determine whether or not you're subscribed to the mailing list.

If you're changing it every time you change your key and you're not changing your mailing address, it will say, "I'm sorry, I won't accept your mail because you're not subscribed." Is there a way to address that at the BATV?

MR. CROCKER: Well, BATV is all about addressing things so there must be. Sorry, but not really. In doing any retroactive change to an infrastructure such as addressing, the likelihood -- where we're

superimposing metasyntax on a string that's had no global syntax to speak of, the dangers of exactly what's just been described are pretty high.

As near as we can tell, this is the only example, which I don't mean to make little of it, but it's sort of astonishing that it seems to be the only case we know of where there is a concern, and the fact that it occupied a fair amount of our time talking about it, and we were going to try to completely accommodate it, and a different approach was taken.

We talked to the Easy LN, Easy MLN folks, and it doesn't sound like it's that big of a deal for them to change. In point of fact that's the major reason for having a standardized metasyntax so globally folks can recognize the syntax, and if they really want a core part of the screen, they can get that and ignore the rest. They don't actually have to understand any of that rest.

MS. DODSON: Can you put it back there and then we'll take your question.

MS. OLSON: Hi, Margaret Olson, O-L-S-O-N. I just wanted to get back to some of the points about costs made by the question of earlier. I think that for small senders, the majority of them are sending through an ISP with a separate domain in the scenario she

outlined, where she had Road Runner and her own domain, so I think in considering the cost of these things, it's very important not to overlook the reality of the many, many small businesses out there and the many small senders who don't really have a technical administrator, don't know what DNS is.

There's quite a bit of infrastructure above the cryptography level that has to go into place in order to implement and deploy this kind of solution, which I acknowledge that from a technical point of view is very complete and very, very attractive, but there's another complete layer that needs to be in place in order for it to be an easy, cost effective thing for the majority of small domains out there, and you're free to comment. I guess that wasn't a question. I apologize.

MR. FENTON: Sure, that's fine. Actually there are some ways that that can be done fairly effectively without expertise on the part of the user. Frequently if these small domains are something that you register as part of your ISP service, your ISP will register the domain for you, they'll operate the domain name system things for you, and basically give you kind of a turnkey domain.

One of the things they could do in the process of doing that, if you wanted to authorize it, would be

to say, all right, from my domain, I would like to have my ISP do the signing for me so you could -- the domain that is operating your DNS just has to authorize its own keys for your domain, and they could either use the same keys as they used for everybody else's mail or maybe for a slightly higher charge and a little bit more security, they would offer to sign your messages with your own key but they would do the signing for you. But they would do the key management for you, and there really isn't anything that you need to do other than ask for the service.

MR. CROCKER: I would like to stress for folks that Margaret's question is just as important as it gets, that we can't get authentication for free, and the different approaches to authentication have some widely varying costs. Some have computing IO costs. Some have administrative costs.

The encryption based ones that we're involved in seem to have relatively modest and relatively stable rather than ongoing administrative costs, but, no, it's not free.

MS. DODSON: We have a question over here.

MR. BOTZER: Bob Botzer, that's B-O-T-Z-E-R with Verfeyes, V-E-R-F-E-Y-E-S, and my question is for Miles and Jim regarding -- I would like you to comment, if you

would, on the adoption of a standard cryptography algorithm to be able to decrypt the messages, and a question for the entire panel regarding -- I don't know whether we're talking about competing standards here or collaborating standards here.

There's a lot of talk about collaborating standards, but in that case how does one tell the difference of which format message and where to go to interpret it properly?

MR. FENTON: Sure. Well, first of all, we weren't actually encrypting the message. I hope that's clear. We're applying a signature to the message, and in both cases, the message headers -- I think this is true for DomainKeys. The message header indicates the algorithm that was used to compute the signature, so it's self describing.

Now, it is a little bit difficult when a new algorithm comes along that we want to use. It's going to take awhile before people are going to have the confidence to use it because everybody that's verifying signatures has to implement the new algorithm before they can successfully verify a signature, so there's this sort of barrier to rapid adoption of a new algorithm because people will want to sign messages with algorithms that everybody can understand, so that was

the first part of the question.

In terms of, I missed part of the second part. It had to do with collaboration?

MR. BOTZER: How do these all fit together or how do they interrelate?

MR. FENTON: Well, I would put what Dave Crocker described BATV being as in a somewhat separate category because it really addresses a separate but very important problem that we have with the handling of bounces. Some domains, people that are -- especially people that are subject to say phishing attacks receive just an unbelievable amount of bounced traffic from the attempts to send these messages to unsuccessful addresses.

And they would like -- it's sort of a good way that they know that they're under attack, but on the other hand, they don't want to have to actually accept all of these messages.

In terms of DomainKeys and Identified Internet Mail, we're really solving basically the same problem. We have both adopted portions of the other, so I would say that we're converging, but since we're here with two different proposals, obviously we haven't converged yet.

MR. LIBBEY: So from my perspective I think we

think of the path to standardization as going through real world testing. John Levine had talked in the outset about the necessity of testing all these different proposals in the real world, and that's why we've deploying DomainKeys with our system today, and once we have deployed and gained this real world experience, we'll know a lot better as to what type of changes need to happen.

MS. DODSON: I guess I have to throw in a plug from the NIST perspective in regard to the cryptographic algorithms. There are some fairly well used identified standards cryptographic algorithms for signatures that they were talking about today. Certainly Arsday and DSS is not used as much, and some work in cryptography has been standardized, so we have one here?

MR. HUTZLER: Can you hear me? Carl Hutzler with America Online, H-U-T-Z-L-E-R.

MS. DODSON: Thank you.

MR. HUTZLER: I would love people to comment on a portion called a pretty name or the display name just quickly, and then the other thing I had was a question foreshadowing the next panel on IP based authentication schemes. David brought up a very good synopsis of why path based approaches do not address all of the different aspects of how the email infrastructure

is being used and how SPF or Sender ID, he alluded to it anyway, may break some of those pieces of the system.

He also alluded to the fact that domain or public private key or encryption based approaches have been tried many times before and have been difficult to implement on a wide scale, although we hope that that will occur in these, and my question is for each group, for each person to comment, should we be looking at IP based path approaches as a positive indicator and not necessarily a negative indicator if those approaches fail or break in some way while we look to cryptographic approaches as sort of the Cadillac solutions.

Maybe this is coming from an engineering perspective. Could SPF, Sender ID approaches flag mail in a positive way or hopefully a large percentage of the mail that already does meet those criteria that does not have the complex paths that it does take?

MR. CROCKER: What do you mean flag in a positive way?

MR. HUTZLER: Perhaps treat that mail as you know that it came from a certain domain, that type of thing.

MR. LIBBEY: The first part of the question was about pretty names and display names so I think all of these proposals validate the actual email address and

don't touch the display name or pretty name, and I'll leave that up to the mail user agent to display as they would like to.

As far as whether path based authentication techniques can be used for positive identification, it's certainly possible. It's definitely a way that these type of proposals can work together. We do think that path based authentication can be used for positive identification, but they have some significant problems in the identification of forgery, and that's where cryptographic solutions would excel.

MR. FENTON: With respect to the pretty name issue, does everyone understand what the pretty name is? It's like a person's name that appears just next to their email address. We've really made an effort to not require changes in mail user agents for initial adoption of Identified Internet Mail. We think that that takes a relatively longer time than it is to just get signing and verification going in the mail servers of some domains.

So as a result of that, we've got a fairly strong recommendation in our specification that if the message is verified as coming from something other than the mail address that would be displayed to the user, that you ought to actually edit the pretty name in order

to make that evident.

It makes a lot of people uncomfortable, and I hear Dave breathing deeply next to me here.

MR. CROCKER: Wait a minute.

MR. FENTON: I'm sorry, I should let you comment for yourself.

MR. CROCKER: I sighed deeply, not heavily. I'm sorry.

MR. FENTON: So we really think it's important to do something, whatever it takes, in order to make the address that was verified visible to the user.

In terms of the issues with deployment of public and private keys, by relying on the domain name system, which is not secured, at least not today, we're kind of making a trade-off against absolute security in the cryptographic sense of what we're proposing versus making this easy to deploy.

So the reason that we do that is because we need to understand what the consequence of a failure of the system is. The consequence of a failure is that mail acts more like it does today so we're really trying to discourage people from using this infrastructure for anything other than decisions about email messages or potentially decisions about other sorts of messages like on instant messages or potentially Voice Over IP in the

future, but we don't want people to have the illusion that there is a completely secure infrastructure because it's not.

MR. CROCKER: I would like to build on what Jim just said because there really is another key to the purpose of this mechanism. It is not trying -- these mechanisms are not trying, for example, to compete with PGP or S-mont. We are not trying to sign messages at any kind of legal level about the content that might be used in essentially a contractual way, and that's another example.

When we are talking about doing any kind of authentication, we need to be very, very clear about what it's used for and what entity is going to use it, and one of the debates that's going on among the technical community working on these is exactly what entity is going to use this.

Now, there is a consistent interest in displaying the information over to the user, so Carl's questions really gets to the heart of that, and as Jim described, there are some approaches that are discussed. Initially I assumed that that was the right thing to do, and I've now come to believe it's actually a very bad mistake.

It's not a mistake to show stuff to the user.

It's a mistake to think you have to. I think these authentication techniques are intended as input to some filtering mechanisms, and they might be in the MUA and they might be in the MTA, and they might be in the user level and they might be in a transfer level, but the primary purpose of these signatures is not for reflecting information to the user, but to provide input into a filtering process.

I think by worrying too much how this gets reflected to the end user in display, we are finding some design distortions that we have to do, and that that's actually making things more complicated.

MR. MATHEW: John Mathew from Obiqua Interactive (phonetic). It's M-A-T-H-E-W. This question/comment is relating to the BATV. I completely agree with the concept and the principles of protecting and verifying all the key components of email.

One of the challenges that still exists today is the treatment of email, even the bounced email back to the large senders and to themselves. Particular x-headers or other types of headers are struck out, so there's no consistent treatment of the bounced email, so in your scenario, that signature may be stripped out by some of the intermediary servers, so how do you handle that?

And just a larger question in terms of making sure that any of these authentication solutions work, there's an underlying assumption that there has to be some consistency in the bounced headers and the messages and leaving certain headers intact.

Is there any kind of effort that's going on today to make sure that bounces are consistent, these headers are consistently included, and if not, one of the efforts or the results of one of the Summits can be that the ISPs get together and make sure there's consistent handling and treatment of those bounce messages. I think that any of the solutions we're talking about will have a greater likelihood of succeeding and working.

MR. CROCKER: So your first question is, is there an effort to make sure these things are handled consistently, the answer is no. Your second point is, well, there should be, and I think the answer is, no, there probably shouldn't, and not that it's not a good idea, but when you have many, many thousands of independent administrations across the globe, the likelihood of getting anybody to make things 100 percent consistent in any kind of timely manner is not very high.

In the case of BATV, we're in luck. We don't

really care about the problem you raise, not because it's not an important problem, but because it has nothing to do with BATV. It turns out BATV puts all the signature information in that bounce address. It's not in any other field, and other than the one example we know of of a mailing list that apparently will break on the syntax we choose, in spite of the fact that it's based on the existing standard, that the relays and even mailing lists will not alter that string.

Now, the question you raised actually is of paramount importance for these two guys, and their specs both deal with it.

MR. QUINLAN: So not to let you run away from it too quickly, this is kind of a follow-up to what was just asked, so each of the different proposals take measures in order to survive inadvertent modification of the message.

I was wondering if the panel could comment, and this is particularly interesting to this group or to the Summit here, about some of the prescribed changes by the path-based systems in order to maintain that path information as you go along. Specifically I'm wondering about incompatibilities of, for example, SRS modifications for SPF, how those could conflict with BATV or header decisions for Sender ID which could

conflict with DK or IIM.

MR. FENTON: It's certainly true if you change the bounce address you've broken any signature on. I don't know of any header addition that's been proposed for Sender ID that would be incompatible with Identified Internet Mail. We can base the signature that we apply on a couple of different header fields. That aspect of the specification is likely to evolve a little bit, but there isn't anything that's fundamentally incompatible there.

MR. LIBBEY: I think the same is true for us.

MS. DODSON: One more question.

MR. ANDERSON: There was a meeting earlier this year, January 20, in Boston where we all absolutely froze to death, but we managed to get I think most of the players that were working on this together in one room, and Meng got up and described SPF and the Microsoft people, Harry got up and described Sender ID, and at that point somebody observed, guys, these things are so much alike, you have got to put them together.

Not doing that will really significantly delay implementation, so I would make the same observation right now, and that is these things are so similar, I don't know what you have to do to get it together, but I think it's absolutely essential that you come up with

one proposal. Dave Anderson.

MR. FENTON: I agree one of the things that's going on right now that leads to that is the experimentation that's going on both with DomainKeys and Identified Internet Mail. We just published an open source implementation of that on Source Forge, and so that will help I think flush things out in terms of what aspects of which proposals are the strengths and really the effectiveness of these proposals I think isn't so much in terms of the number of messages people get signed. It's the number of messages that verify in all the different use cases. That's what we need to find out with the experiments.

MR. LIBBEY: We absolutely agree. Particularly the real world experience is going to tell us a lot. We don't want to make the same mistakes that happened in MARID, and without that real world experience, so that's why we're focusing on getting deployments out.

MS. DODSON: I would like to thank all the panelists. I think you've all done an excellent job.

MS. DODSON: I appreciate all the good questions too from the audience. There is a one hour lunch scheduled, and if you all look in your packet, there is a "Where to Eat in the Vicinity," but the Summit will start again at 1:30, is that correct?

MS. COLEMAN: That's right, Donna.

MS. DODSON: So everybody needs to be back by 1:30. Thank you.

(Applause.)

(Break in the proceedings from 12:30 to 1:30 p.m.)

AFTERNOON SESSION

(Resumed at 1:30 p.m.)

PANEL 3: EMAIL AUTHENTICATION PROPOSALS: IP/DOMAIN BASED APPROACHES

MODERATOR: WILLIAM E. BURR, NIST

PANEL MEMBERS:

HARRY KATZ, Microsoft

DOUGLAS OTIS, Mail Abuse Prevention System

MENG WENG WONG, Pobox.com

DAVE CROCKER, Brandenburg InternetWorking

MR. BURR: Folks, can I ask you to come in and take your seat so we can get the session started and we can stay on time? I'm Bill Burr from NIST, and like my colleague, Donna, I would like to express our pleasure at being invited to participate in this, what's turning out to be very interesting and productive workshop, and I would like to thank Donna Dodson and all the FTC crew that did 99.99 percent of the work to put this together.

We've learned about cryptography in the last session, and the group that I work with at NIST actually deals with cryptographic standards, so in a way I think we ought to quit while we're ahead, but obviously there's another side to this, and one of the things that I've learned in my experience with PKI is you say the

word "PKI", and 90 percent of the world winces.

So for this session we've actually got two speakers, and we're going to have I think a total of four panelists, which is we're going to have Harry Katz from Microsoft give a little presentation on Sender ID, Doug Otis from Mail Abuse Prevention Systems talk about CSV, and I think they're sort of opposite poles of the same spectrum as it were.

And then Meng Weng Wong will be available also to answer questions on SPF, which is I think now widely subsumed in the Sender ID, and Dave Crocker is here to provide further balance.

So we have already heard a lot about SPF, and so it might -- Sender ID, excuse me. We've also heard about SPF, and it might as well have been well to almost have started with this discussion.

What I would like to do actually is, I won't say we've exhausted the topic, but we spent plenty of time talking about licenses at this point, so I would like in this session not to revisit licenses particularly, but to go on to talk about SPF, what it is and what it would do if we did it.

So with that I would like to introduce Harry Katz from Microsoft who will speak about SPF or Sender ID. I'll get it right yet.

.

(Applause.)

MR. KATZ: Well, thank you very much, and good morning, everyone. My name is Harry Katz, as I was just introduced, and I would like to ask you all to please pretend it's still morning, you haven't just had lunch, and you don't really feel like the need to dose off for a little siesta right now. Pretend this isn't the Bermuda Triangle of the afternoon. Instead pretend that you just had your morning Espresso or your Grande double non-fat, no-foam double cup latte or whatever it is that you order, and we'll try to keep this as lively as I can.

I want to just give a brief overview of what Microsoft has been doing and what our anti-spam strategy is, talk about a little context of why we think we need email authentication, and talk in much more depth about the Sender ID framework and then look at the implementation considerations and the benefits of the Sender ID proposal.

So at Microsoft, spam is one of our customers' number 1 complaints, particularly about email, and we've been working on the spam problem with a great deal of focus for at least the last two years, and we've kind of gone through a five prong or five pillared approach that

consists of technological innovation, industry collaboration, strong legislation, support for the CAN-SPAM Act, strong enforcement of that legislation and consumer education.

We've been very active on all five of those fronts, and clearly the Sender ID proposal is something that fits into the technological innovation aspects of that strategy.

We think it's important because it does add this dimension of email authentication to the whole question of spam filtering. This slide is an attempt to answer this question why we think email authentication is important. Over the last I would say two years, a great deal of the focus and the investment in anti-spam filtering has dealt with content filtering, trying to identify whether or not the content of a message is good, bad or ugly based on the analysis of the actual message content.

I think we've made tremendous progress as an industry, as a company too, but as an industry, we've made great progress here in terms of increasing the effectiveness of those content filters. I would say there are many products on the market today, not just from my company, that can give you filtering success rates of around 90 percent in terms of the catching the

spam that's coming in.

There are problems that remain. There's still obviously some spam that comes through and, we can't crank up the aggressiveness of those spam filters without risking increased number of false positives, that is to say legitimate mail that is misclassified as spam.

So we need to move forward now and take additional steps to just -- in addition to rather just looking at the content of the message. We need to take a look at who is the message from, who is the sender of the message and see if we can make some determination about the likelihood of mail from that sender being good or bad, and this leads us to the notion of sender reputation systems.

Now, these have been around for awhile, and in their initial form they take the form of IP reputation systems, and these are well known as the various blocklist services that are out there today and fairly widely used, and as well we're starting to see some IP based solutions that list good senders as well.

But as I think it was Miles Libbey who pointed out in a crypto presentation, IP based reputation has some problems because organizations can share IPs with other organizations. Also many companies, large companies in

particular are constantly bringing up and taking down servers so IP addresses change.

IP addresses change, and that means that you have to start all over in terms of building up a reputation for a particular IP address, so it's much better or much more resilient to those kinds of changes if you can hang the reputation on the domain rather than IP address.

In addition, once we have some notion of where domain message really originated from, we can give feedback to the originator of that message to help them improve their behavior. We've already seen a lot of movement on the part of mail senders, particularly ISPs and ESPs, over the last year adopting best practices, like Port25 blocking and rate limiting.

You're going to hear me talking about publishing SPF records. There's digital signatures as other things that senders can do, proof of work ideas a sender can do to distinguish their mail from spam, but all of that hangs on this notion that we can identify with some accuracy who the sender is, and in particular what the domain is, and that's where Sender ID comes in.

So what is the Sender ID Framework? I'll be bold enough and call it an emerging standard, and it's a merger of a number of proposals and some feedback that

we've received from various quarters, in particular it incorporates the sender policy framework that was first written up by Meng Wong and his partner, Mark Lesner, and a great many others who contributed. I know Hadmut Danisch is in the audience. He's one of the progenitors of this whole idea as well, and it also emerges in a Microsoft Caller ID proposal that was being developed by Microsoft internally around roughly the same time as SPF.

Both these proposals got submitted to the IETF MARID working group and we benefitted from the feedback of that working group, and so the document and the specifications that are available today reflect the merger of those proposals and all the feedback.

Along the way we've been coordinating and consulting with a number of organizations, stakeholder groups within the email community, and we're gratified to have feedback and support from a large number of organizations.

Now, when you're looking at a problem like this where you have a mail system that has been deployed across the planet over the course of 20 to 25 years, where it's in use by somewhere between half a billion and a billion people worldwide, you really have to be very careful about what you do and how you slice the

problem, and so this slide is an attempt to capture some of the trade-offs and design decisions that we've been making.

Now, it's certainly possible to choose other sets of trade-offs and other parameters, but this is where we think sort of the balance needs to lie for Sender ID at any rate. We think it's important to give domains the ability to protect their brands and their domain names.

We also think it's important to be able to hold those domains to account for the mail they send. I mentioned the scale of the Internet so we need to ensure that the system can, in fact, be deployed at Internet scale and can he easily adopted, and that's not to say that this is a silver bullet or that it's going to be totally painless or totally free or we're going to solve all the problems at once. We're trying to take a measured and reasonable approach to solving a significant piece of the problem.

So the Sender ID framework now is really composed of four elements that you see here. The first is what's called the SPF record, and I think you've heard some mention of this earlier this morning. This is the record that we request organizations, sending organizations to publish in the DNS, in the domain name

system, the global Internet directory that identifies the authorized outbound email servers for a domain.

Once an organization has published that record, then receiving organizations who get mail from that domain are now able to perform one or both of two different checks or two different validations, one of which is a validation of the Mail From address, and another which is a validation of what we call the purported responsible address or the PRA. So either or both of these two checks can be implemented on the receiving side.

In addition to that there's an optimization or a minor enhancement to the SMTP protocol itself to allow the purported responsible address to be sent with a message envelope so that validation of the PRA address can occur earlier in the message processing cycle, so those are, if you will, the specification elements of the Sender ID framework.

So how does Sender ID work? Well, the first step in this awesomely animated graphic is that organizations publish in the DNS their outbound -- the IP addresses of their authorized outbound email servers. Then they just send mail as normal, and at the receiving end organizations decide which of the checks they're going to perform.

They isolate the appropriate domain name, make a query to the DNS system to look up the SPF record for that domain, and then they try to do a match. They're looking for match on IP address. Is the IP address over which the specific message was received -- is that IP address authorized as one of the official outbound email servers of the domain?

If it is authorized, then there's good evidence that the message as originated properly from the domain it claims to come from. If it's not, if there's no match, then you have some pretty good evidence of spoofing.

I want to talk for a minute about the two checks, the PRA and Mail From Check, to sort of compare and contrast these a little bit. First of all, the Mail From check is based on what is known as the bounce address or the RFC 2821 mail from protocol address, and by contrast, the purported responsible address is actually derived from the message headers.

We tried to look through the headers of the message to identify and isolate the identity that's most likely to be responsible for injecting the message into the mail system. We think one of the advantages of that it is more likely to perform a validation on an email address that is ultimately displayed to the user when

they open the message.

Now, at Microsoft we're the ones driving the PRA check, the original authors of SPF. We've driving the Mail From check. We've now sort of essentially merged them under this umbrella of the Sender ID framework. I should say there are some advantages and disadvantages to both systems, and I would also say they're focused on different parts of the problem.

The Mail From check I think is at least originally as it was conceived seems to be focused on solving the false bounce problem or the joe-job problem. Dave Crocker described this a little bit earlier in his presentation where an attacker sends spam. It's spoofed, and all of the non delivery reports and other notices get sent to some innocent victim.

From the perspective of the PRA, we think because this is focused on validating an identity that is available and displayed to an end user in most cases, that this is something that helps us to start to address the phishing problem, so these things are we think relatively complementary but nonetheless focused on different aspects, different takes on what the problem is.

Now, once you've performed a Sender ID check, you get a result back from that exercise, and you have

the choice of certain actions to take on the basis of that, on the basis of that result. You could accept the message as good. You could reject it outright, if you so choose, or more likely, and this is certainly the path that Microsoft will be pursuing and I know that the Hotmail folks are pursuing in their implementation, they will simply use the result of the check as an additional input into their filtering decision.

Now, we can expect over time that as adoption gets broader and more and more people are publishing SPF records and more and more receivers are validating, that the weight of the Sender ID check will increase in these filters, and so it will be increasingly important for organizations to publish those records, assuming that thing is adopted and succeeds, but at least in the initial stages, I certainly wouldn't recommend outright deletion or rejection of the messages based on the results of the Sender ID test.

It's important we think, perhaps not critical in the early stages, but over time we think it's going to be important that when we do some of this validation, that we convey to the end user of the message which identity was validated and whether or not that validation succeeded, so that's not another thing, another action, if you will, that can take place on the

basis of the check.

And just to reiterate the point, Sender ID is a proposal that tells you something about the sender. It tells you nothing about the content of message per se.

So it is perfectly possible for a spammer to go out and register their own domain name, publish an SPF record and send you spam which passes the Sender ID check. In fact, I think Cipher Trust, an organization in this space, published a study a couple weeks ago citing that a large number of spam actually passed the Sender ID check. Frankly I think that's fantastic news, and to me it's proof that this is going to work.

If we get spammers registering their domain names and publishing SPF records, they're effectively stepping out in the open and saying, "Here I am, shoot me," and that's what we want.

Now, I've given this presentation on quite a number of occasions, and there are a number of people in this room who have had this inflicted on them several times. In fact, last week I was at a meeting with Jim Fenton who's at Cisco and made the point that this whole email authentication effort is beginning to resemble World Cup skiing, and it's like there's this cluster of athletes that all know each other, and sometimes they're competitors, but off hours they're friends, and they go

around from place to place and they do their thing.

Well, we're doing that here, in perhaps not quite so exotic surroundings, but there's great opportunity for cooperation and collaboration, which is great, but as I said, I've given this presentations on a number of occasions, and I always get two kinds of feedback.

The first says there's not enough technical detail in my presentation, and the second feedback says there's too much technical detail, so a fair warning, the next few slides are going to be the technical part of the presentation, so pay attention. There will be a quiz at the end, and if you don't pass, then you will have to go to the Inbox Conference in Atlanta next week and listen to me give this talk all over again.

Okay. So I want to talk a little bit about what these SPF records are. We've been telling everyone you need to go out and publish these things. They're records that indicate various policies, if you will, about the domain that has published them. The first record -- I won't go into detail on all these, but the first record is really sort of the base case, and this is one where a domain says, hey, we never send mail, this is a domain name that is registered for other purposes, we never send mail, and we only have version

tag and this minus all indicator at the end of the word. If you received mail from us, we don't send mail so it's spoofed.

The next example shows you how a domain that has -- typically a small domain that may only have one or two mail servers that are doing both inbound and outbound processing. There's this little key word in there called MX. That basically says go and look at our DNS MX records, those are the mail exchanger records that tell you what the IP address of an inbound mail server is. Those are also valid as our outbound mail server.

I'll skip down a few. Is the fourth one here is one that allows an organization to designate a third party or perhaps a parent domain or a subdomain as being authorized to also send mail on behalf of the domain, so it's sort of an out-sourced scenario where you can say, Hey, these are my authorized outbound email servers, but in addition go and look at that domain's SPF record and their authorized mail servers are also okay for our domain.

Now, there are a number of scenarios and delivery paths as messages travel, as they go from ultimate sender, in this case Alice@example.com to the receiver, Bob@woodgrove.com. The more straight forward

case of course is mail direct delivery, but you can also have situations where there are intermediaries, what we call agents in between along the message path.

Some of those agents act on behalf of the sender. Some of them act on behalf of the receiver. Mail agents that act on behalf of the sender such as list servers which distribute mail to all of the subscribers of the list, to mobile carrier networks that send mail on behalf of the user of a little mobile device out to the Internet and guest email services, like electronic greeting cards and electronic invitations, emailing this newspaper article to a friend, et cetera, that send mail on your behalf when you don't really have an account on their network.

Forwarding, the quick example of an agent that acts on behalf of the receiver, so I want to look at a couple of these scenarios and sort of explore what it is that the senders need to do in order to be compliant with Sender ID.

For direct delivery all they need to do is publish their outbound server records in DNS. In other words, they just need to publish their SPF records and they're done. That's it. This is the really geeky part. This is sort of a transcription of what the SMTP sessions looks like. In the interest of time, I'm not

going to go over this in any kind of detail, although as a technologist this is the part that really excites me, but I will only point out here that in this particular case of direct delivery, the Mail From address in the envelope and the From address in the body of the message are identical.

So in this case it really doesn't matter whether you're doing a Mail From check or a PRA check. You're both checking the same domain.

Now, in the case of mailing lists, as I mentioned earlier, they fan out mail to all the members of the list. What they need to do in order to become compliant are two things. One, publish their SPF records and two, they need to ensure that there is some identification of the mailing list server itself or the mailing list domain itself in the message, and the vast majority of the mailing lists do this today already.

They use a list owner style of address, and they use this in the Mail From command, and many of them also insert a sender header in the message, so most mailing list senders, not all, but most of them are already compliant today. All they need to do is publish their SPF records.

For forwarders, again in this case we've got the classic example of a college alumni account so Bob here

has an Alma Mater.edu and he set it up to forward mail to himself at his office at Wood Grove. What does that forwarder need to do? Well, I'm sounding like a broken record, but they need to publish their SPF records, and they also need to indicate at some point in the message what the domain of the forwarder is. They need to identify themselves as having, if you will, touched the message.

One way to do this is to insert a Resent From header into the message. There are also schemes that they could use to rewrite the Mail From address.

All right. So I wanted to briefly touch on some of the implementation considerations here in terms of what some of the costs are, what people need to do in order to become compliant with Sender ID. They're really divided the ecosystem into three broad components: Senders, receivers and these intermediaries like forwarders and listservs.

What senders need to do, the main costs for senders is to fundamentally track down what those outbound IP addresses are, get their SPF records published, and for large organizations to maintain those records, and we need to recognize that that is an ongoing administrative cost.

It's also important for organizations, large or

small, that are out-sourcing their email services that they contact those out-source providers, make sure that those guys are publishing SPF records and make sure that they have the necessary directives in their SPF records so that the messages that emanate from those out-source providers are seen as legitimate.

Receivers in the short term, we would obviously want them to upgrade. There's no software upgrade required for them to perform either the PRA or Mail From check, in a little bit longer term, changes presumably to clients to display some information about the results of that validation.

The email intermediaries like list servers and forwarders, they're a sender like everybody else, so they have to publish their SPF records, and they also have to probably make some software changes, if they haven't done so already, to indicate that an address under their administrative control has taken responsibility for introducing the message on that next hop.

You heard this morning a panel on the cryptographic approaches. I just wanted to take a brief minute to compare and contrast these two approaches. We think they're complementary. There are some strengths and weaknesses in both. Neither of them are going to

solve all the problems, but we do think that there is a great chance that they can reinforce each other.

Sender ID is something that validates the last hop of a message. If a message transmits through several of these intermediaries, Sender ID only validates the last of those jumps, the last of those hops. The cryptographic approaches bring the promise of validating the original author or the original sending domain of the message. The proviso is that the digital signature must survive transit through the entire passage path.

The Sender ID proposal is designed to validate the domain of the sending organization. The crypto solutions are also designed to validate the domain but also have the potential to also be used in user based validation.

There's a difference in deployment. For senders deploying Sender ID today, all they need to do is publish their SPF records. There is no software upgrades required, so we think that's a great advantage, and maybe that's one reason why we think Sender ID -- these IP based solutions is something that can be deployed quickly and right away.

With cryptographic based solutions, you need to have software at both ends before you get the benefits.

You need to have the senders who are actually creating the signatures and the receivers who are validating them.

Both systems tell you something about the sender of the message, and so have some vulnerability to certain kinds of attacks, and therefore both systems serve as inputs into further reputation systems that are based on the sending domain, so we've been in discussions with Yahoo! and Cisco and a number of other folks talking about these cryptographic based solutions. We look forward to seeing these continue to evolve, and we think they're complementary with Sender ID and the IP based approaches.

I just wanted to quickly wrap up now with an overview of what I think the benefits of Sender ID are. First of all, it provides the ability for senders right now to take immediate steps to protect their domain names and their brand names against spoofing and phishing attacks. We think it's amenable to rapid adoption in terms of simply deploying the records and not having senders at least required to upgrade their software right away.

It's a basis for reputation and accreditation systems. It's a basis for reliable use of safe lists that are built on the domain name of the sending

organization. Receivers get the ability to now validate that the sending domain is in fact who it claims to be, and what that does is give us additional input into the spam filtering decision, allows us to crank up the aggressiveness and rigors of our spam solution, with reduced risk of false positives.

Finally this is an opportunity and I suppose a challenge as well for the industry to come together and collaborate on solutions. All of the anti-spam solutions that have been created thus far are themes that corporation organizations can unilaterally develop and deploy. You can go out and buy or select a whole host of spam filtering software, subscribe to an IP block list as you choose.

Sender ID and like solutions are really the first kind of solution that require systematic change to the email infrastructure, and that requires a great deal of collaboration which is a long and sometimes slower process than we like, but it's certainly a very important exercise for us all to go through.

In summary in case you haven't gotten the message, publish your SPF records. Microsoft is going to be starting, checking, doing the validation through Hotmail by the end of this year. I know a number of other organizations are going to be doing the same, and

talk to your MTA providers about getting their software upgraded to perform the Sender ID checks.

So again I want to thank the FTC for giving us the opportunity to come here and present on Sender ID.

Thank you.

(Applause.)

MR. BURR: Our next speaker is Douglas Otis, and he's going to talk about CSV and probably has a somewhat different view of a number of things.

MR. OTIS: Hello. I'm Douglas Otis. I've been working with MAPS for a few years and learning an interesting aspect of dealing with email. I'm not really what you call a professional key class public instructor. I'm more of a geek. I'm going to sound like a geek.

Anyway, are the topics I'm going to be discussion. I plan to walk you through reasons why we need to develop an accurate and lightweight email authentication standard, why security is so key and why some proposals will put us at greater risk, who should be the entity who's held accountable and how to assess their reputation, how problems are addressed with client SMTP validation or CSV, and how the CSV solution will reduce the levels of abuse while also avoiding the security risks present in some of the other proposals.

Before I delve into these issues, here are some of the general terms related to email authentication. Although the source for identification varies between different proposals, the basics remain consistent, and this is important because essentially you're deciding by this identification what ox is going to get gored, because you don't stop spam without reputation or accreditation, and eventually someone gets hurt.

So this identification is not a trivial task in deciding, so identification, who does this purport to be, but in addition which field are you looking at?

Authentication, is it really them? Authorization, what are they allowed to do, and accreditation and also reputation, are they recognized?

Authentication and authorization validates the identity, and that completes the first phase of CSV. Accreditation would be the second phase. In addition, I'll be using a term called mailbox domain. This refers to the owner of the domain to the right of the "at" symbol in the email address, and host domain refers to the owner of the domain operating the SMTP client. That's a big distinction between what Sender ID, FBV and CSV are about. We pay no attention to the mailbox domain.

We may focus on the various techniques used in abusive email, but defeating security remains the

principal method for circumventing otherwise effective spam protection. A system may be compromised, often unbeknownst to the owners, I'm sorry. Where frequently this happens is a way to commandeer and unblock addresses.

When considering email authentication, the identity that needs to be validated is that of the entity ensuring security. This identifier must be relatively strong. Thus this requires direct authentication to ensure the integrity of the system. This entity is revealed by the IP address or the host domain.

It's only the administrator of this address or domain that is able to take immediate action should abuse be detected. The HELO domain is the only name identifier within an email message that can fulfill this role.

Once the administrator has been determined, reputation of this entity is then judged by the action taken upon notice of abuse. In other words, we don't trust IP. IP we view as kind of like the garden gate leading into the front door. The front door should be guarded by cryptographic technologies like Identified Internet Mail or Yahoo! DomainKeys, but that garden gate is important because otherwise the pathway to that front

door would be trampled. So we don't trust it very much, but it has to be there.

The resulting reputation offers protection against a growing torrent of abusive email. Reputation services such as blocking lists base the acceptance of email upon the IP address of the SMTP client, and early reputation assessment of IP address within SMTP session conserves both systems and network resources.

Being early in the session is a critical aspect for email protection schemes. The expense required to keep address based information current, however, with the related difficulties of determining the administrator could be reduced by adoption of name based information.

A name based reputation system will also extend protection to other aspects of email such as email signature systems. Ensuring the name relating to the entity accountable for security of the system is possible by validating the HELO domain. Also a HELO domain assessment can also be done early in the SMTP session.

Its authentication, unfortunately, must be allowed to fail as the protocol now stands. Security's ongoing challenge, whether for a large network provider or grandma's desktop, recipient educated script is found

within HTML messages, which is the basis for enticing interactive multi media, represents a major component of the security threat.

As evidenced by the recent security peril from displaying a JPEG picture, even the simplest script adds risk, unlike a browser where scripts are obtained and executed at the behest of the recipient, email allows scripts to be distributed without recipient intervention.

As a result, the script related vulnerability within email is far more serious due to the ease by which malicious scripts spread. Who should be accountable?

There's a variance granted in RFC 2821 to accommodate a DNS address resource record where addresses drop off the end of the response. This hinders any assurance that all necessary addresses will be returned to ensure the authentication of the HELO domain. CSV solves this issue by utilizing a service or SRV resource record to establish new expectations.

By validating the HELO domain rather than just using just an IP address, a name can be used to establish a reputation of those accountable for security in the administration of the SMTP mail transfer agent or MTA.

The HELO domain parameter is already exchanged by SMTP. Basing reputation on this entity rather than the IP address places accountability on the same entity and does not alter the current email paradigm. Sorry.

Now I'm too far. For some of the new email schemes being proposed, the entity that receives the reputation could be a mailbox domain based on Mail From sender or the recent series of headers within a message. With the new decision, you don't even know when you publish the record which field you're authorizing.

These new mailbox domains authorize SMTP clients through a set of DNS published scripts that describe the mail channel with a comprehensive address lists. Examples of such schemes would be SPF or Sender ID.

The mailbox domain cannot be directly authenticated using an IP address, but would receive a reputation irrespective of the domain's ability to take corrective action if no other domain is offered.

Since security accountability is not encompassed by the mailbox domain address list schemes, litigation may be required to ascertain a negligent party with respect to security and to resolve any resulting reputation issues.

Organizations forced to use a mailbox domain

address list scheme may suffer lost messages or become blocked by a reputation service when security is neglected by one of its service providers that remains unidentified by such a scheme.

Is the mailbox domain reputation bad due to the out sourced customer support or was it their advertising agency that had the security problem? As security is assumed by these mailbox domain address list schemes, the mailbox domain, which often serves as a type of trademark, may be damaged beyond the owner's control. Even going to a different provider will not offer relief because it is the mailbox domain that receives the bad reputation.

The problem of accountability based upon the mailbox domain address list authorization is even more difficult when exceptions are permitted. Such exceptions are enabled by declaring the address list to be open ended. The purpose of this is to overcome issues related to the use of forwarding or the use of kiosk style network access.

Such domains with open ended address lists which assure messages are not rejected -- I'm sorry, should domains with open ended address lists which assure messages are not rejected have their name tarnished when their mailbox domain becomes exploited. There are some

proponents that say yes.

Added to the problems defending the reputation of a mailbox domain, there's a lack of agreement as well as intellectual property issues resolving which mailbox domain is checked for authorization. SMTP is not end to end. email travels through several separately administered systems before arriving at the ultimate destination. These multiple administrative regions make spoofing and mailbox domain difficult to prevent when each region may have checked different headers. The mailbox domain selected by these authorization algorithms may also be invisible to the recipient.

Without consistent checks within the email channel, there can be no authorization assurance or accurate reputation assessments made based upon the mailbox domain even assuming perfect security. To make this problem worse, there are many practices aimed at improving security that merge mailbox domains into a common mail channel. Forcing mail to run through the providers's SMTP server used to monitor air logs as a method to discover and exclude abusive customers, but at the same time severely weakens any assurance that a mailbox domain as indeed authorizing the sending of a particular message, nevertheless, using a name that's desired.

Name based reputation in addition to reducing the expense of attracting abusers would be helpful in protecting signature systems that actually authenticate the original source of mail such as Cisco's Identified Internet Mail or Yahoo!'s DomainKeys.

Although these schemes authenticate a name, the name can still be that of a spammer. In addition, method signatures require processing the entire message and offer no resource relief. The use of a name can also override the results of an address blocking list, allowing the owner to change addresses and still retain the reputation.

For an analogy of a fair reputation model, view the mailbox domain as an insurance company. View the SMTP transfer agent or MTA as an insurance broker or advantage and view the mail recipients as clientele.

The insurance broker has an fiduciary responsibility to ensure secure transactions in a timely manner. The insurance broker's reputation is based upon their ability to resolve problems and their offering of only reputable insurance companies.

The insurance broker is identified with the unique name by their license. Clientele are protected by confirming the name of the insurance broker with the insurance company or with the reputation service.

Should there be fraud, transaction logs of the insurance broker are a principal instrument for enforcement. Reputation becomes the principal instrument for consumer protection, perhaps through the loss of the broker's license.

The CSV scheme follows this insurance industry structure. Unlike a mailbox domain address list authorization scheme, CSV validates a unique name rather than offering just a nebulous address for the specific MTA. If there is fraud, it is the validated name of the MTA that's held accountable. The logs of the MTA can be discovered for enforcement purposes, and the party responsible for security and resolving issues is appropriately attributed for any possible abuse.

In this structure the MTA vets the mailbox domain on behalf of the recipient and the MTA's reputation depends on its ability to do so for doing so. In this scheme, the mailbox domain is not harmed by the negligent administration of an MTA, and the mailbox domain and the HELO domain are different entities.

Should there be a problem, the owner of the mailbox domain can freely seek a new provider. This protection is not provided by a mailbox domain address list authorization scheme.

The only name provided by these schemes is that of the mailbox domain and this name will likely be attributed for any abuse regardless of the entity's accountability for security. CSV in any case ensures a reputational service accurately assesses the specific source of any problems and thus allows for the most expedient resolution.

Unlike a mailbox domain address list authorization scheme, there's never a doubt within a chain of transactions which entity is accountable for ensuring security at each step. This entity validated by CSV can also be presented to a filter as a relatively strong mark to prevent spoofing or phishing.

For financial institutions the consumer could also be further protected by publishing a simple name list with HELO domains to make sure phishing attempts are thwarted. Validation of email must concentrate on identifying those able to take corrective action. In general this would be the administrator of a specific host running the MTA readily identified by the host name.

If it were not for a minor defect in SMTP, this code name could be found by validating the HELO domain provided at the beginning of an SMTP session. By repairing this defect, the host name would determine the

entity able to take corrective action as well as the location of transaction logs needed to trace criminal activity.

The CSV, CSA, SRV record, this is geek, I'm sorry, is essential but a simple element needed to repair SMTP. Any complexity regarding the SRV record would have been in respect to implementing a load distribution normally required for this record.

However, the use of the SRV record to authenticate and authorize the client does not deal with this complexity at all. The priority and weight fields intended for load balancing are redefined when used to validate the client. This approach could be used with other protocols as well.

RFC 2821 requires that a failure to authenticate the HELO domain does not cause the session to be refused. This failure occurs when a sizeable constrained DNS elects to drop IP addresses. The address dropped, however, could be the address currently in use which presents a confirmation needed for authentication.

In the normal use of DNS, such a reduced set of addresses still locates the related server and is not a problem. The address lists are often created in random or round robin fashion, but this ordering technique also

serves as a crude form of load balancing with a dropped address is varied per request after the expiration of these records and the local cache.

CSV revolves this issue by utilizing a service resource record to establish an expectation that all possible addresses for the SMTP client will be present. This record type was engineered to return a set of addresses for a service where the client is expected to implement more elaborate load balancing.

The use of the SRV record does not require the double entry of addresses needed in address list scripts as address generation is an automated function of this record type. This record type simplifies maintenance without incurring additional lookup overhead.

The service resource record was introduced in 1996 and adopted by Microsoft when they transitioned from DNS from their service. Additional fields within the service record also permits the domain administrator to assert various mail policies beyond what would have been possible just using the address resource record.

CSV currently uses these fields to specifically authorize a host for sending mail and to note the current version of the record. There's an on going discussion about potentially defining the use of another bid in this field. CSV validation of the SMTP client

can actually be achieved with less overhead than that incurred today. The simplicity of CSV where just a record is added for the outbound SMTP client should not detract from the significant impact of this change.

Making the HELO -- I'm sorry, making the validated HELO domain visible to the filter should offer the strongest form of anti-spoofing protection possible without the use of signatures. To enable the detection of messages with possibly spoofed mailbox domains, CSV permits the mailbox -- I'm sorry, it permits the mail channel information to be published without incurring excessive maintenance for the provider or risk for the recipient.

Once the HELO domain is validated constraining a mail box domain to a root name list with the HELO domains is a protection mechanism for financial institutions can be achieved within a single DNS lookup. As an option used with CSV, this -- I'm sorry.

Such a name list can take advantage of the DNS built in name compression and be assured to fit within the DNS lookup. In addition the association of these two entities can be open ended without inviting an exploit because a mailbox domain would not be used to establish reputation.

On the other hand, a comprehensive address list

of the mail channel defined with scripts may require hundreds of such lookups for every message.

The only name ensured from the address list approach is the mailbox domain. As a result these address list schemes run a much greater risk of misapplied reputation. In addition the existing mechanism is ideal for a criminal sending from a compromised system as a means to obfuscate the range of addresses they're claiming. CSV however uses the native records currently available within DNS, the nationally constrained range of addresses that can be claimed.

The implementation of the mailbox domain address list schemes require one to ten DNS text resource records containing scripts to be parsed by the recipient. The sequential nature of this parsing from several DNS servers is ideal for a cache poisoning exploit.

Often an operating system utilizes many ports to multiplex communications between program threats. Normally this is not a problem as a DNS lookup would be to a single name server and thus would not expose the port employed by the system.

In the process of parsing the scripts, however, a miscreant would only need to place the nefarious email server before the name server they wish to

override with poison records.

The sequence of lookups to different name servers exposes the port in play, and a single script can make more than 300 replicate requests within an equal spoofing of 300 DNS responses which has a 50 percent probability of poisoning the cache. Such a poisoning scenario can be easily done using just the DNS connection.

These DNS published scripts allow additional risks. The overhead needed to resolve potentially hundreds of DNS records specified by the scripts can easily overwhelm TCP network traffic with predominantly UDP traffic.

To make this worse, to guard against denial of service attack, both schemes have elected to ignore UDP exponential back off and simply failed to lookup prematurely. This lack of congestion avoidance is a common but dangerous area in new protocols. That is, that if there's a requirement that scripts without changing revision can be extended.

There's no way to predict the eventual size or complexity of the script. It is clear from its onset that promoters of these scripting schemes ignored advice provided by the DNS working group.

The address list needed by the mailbox domain

schemes overwhelms the design scale of DNS by requiring a comprehensive set of addresses for all hosts that may send mail for a particular email domain. DNS was designed primarily to provide a small address list for a specific host. CSV stays within these constraints.

In conclusion finally, security is not a solved issue, nor will security be fully solved any time in the near future. The reputation service must assist in identifying compromised security. The reputation server and the email service provider must work closely together to guard the email system.

In preparing the HELO domain authentication, using the record has a benefit of also requiring specific authorization by the administrator. Compromised systems would only be enabled by cooperative name servers and thereby would increase their exposure from such an activity.

CSV does not represent anywhere near the same risks by those imposed by systems that put active content into DNS. CSV is simple to implement and does not require any sequential lookup or the parsing of scripts.

By ensuring reputation as asserted on the host domain, those accountable for security are tracked by the reputation service. CSV does not alter the SMTP

protocol currently and permits the same freedoms currently enjoyed.

For exigent situations, CSV also allows the mailbox domain to be safely constrained to a prescribed mail channel without creating additional security risk. email authentication is about security.

Thank you.

(Applause.)

MR. BURR: Okay. Is Meng Weng Wong on the room now? Well, I keep trying. If he would like to participate in this panel, it's time now. I've been told he was wearing a cape.

While we're waiting, I would like to ask a question, and then people counter -- Mr. Weng, would you like to join us up here? Mr. Wong rather. All right. I have to collect myself here now.

I would like to ask people if either of these systems that we're talking about here are more than an expedient to get something in effect quicker than we can put a cryptographic solution in place, or if they have a long term purpose in the scheme of things.

So, Douglas, you start.

MR. OTIS: Well, in terms of providing a lightweight security mechanism or at least a way of knocking down the majority of what you have coming into

your mail system, I think there is something that's needed to kind of ferret out the majority or the bulk of what you're going to be processing for your email.

None of the very secure systems using signatures offer any relief in terms of network resources or system resources, and essentially the IP Gateway, if you will, does offer the garden gate kind of protection that protect the pathway to the front door, and I think that that's going to be a long-term requirement.

It's not something that's going to go away, but it's something that you can't really rely on. People can step over it rather easily, and so you have to understand that the security there is very weak. The authentication must be as direct as possible, and I think that's something that we're going to need for a long time to come, and that's why I think it's important to fix that little blemish, if you will, in SMTP.

MR. BURR: Anybody else want to hack at that?

MR. KATZ: Well, as I said in my presentation, I think we believe that the IP based authentication can be complementary or is complementary to signing so I do think there is a long-term for both of them.

MR. BURR: Anywhere else? If not then.

MR. CROCKER: My view is that there is a need for information about the operator which is the MTA, and

information about the author or the sender, and as Harry says, this is quite complimentary. The means of providing that information is an open area of research that we've got people exploring, so whether it's using some form of IP authentication or encryption authentication is some of what we need to try to understand better.

MR. BURR: Okay. I would like to throw it open to the floor, and I would like to ask people to use microphones and to make sure and state your name, so down here.

MS. ROBBINS: Bill, we have one question on a card. Maybe I'll read that one first, and then I'll walk over there. This question is for Harry:

"Doug Otis has stated that CSV's authentication of the HELO domain has numerous benefits over authentication of the carry or mail from. Could you comment on this?"

MR. KATZ: I won't go into much detail on this. Let me say at the outset that I guess I would have to say I don't have any strenuous objections to the CSV proposal, and I think that authenticating the HELO domain or the HELO domain is a fine thing to do.

My view on it frankly is it just doesn't give you enough of a benefit to justify the cost. I think

that the administrative costs of CSV are roughly comparable of that to Sender ID in terms of the amount of information that gets published, and I think that Sender ID goes a little bit farther in terms of providing information directly about the domain that is contained in the message and allows us to take some further steps in dealing with the phishing problem.

MR. BURR: Doug, do you want a piece of that?

MR. OTIS: In terms of reputation, there is virtually no value in the mailbox domain that you might obtain from anything that might be authorized by Sender ID. The problem with that is essentially hearsay.

We spent a fair amount of our effort in not only providing the reputation services, but we have an equal amount of effort in providing discovery that goes along with that, and so we're turning the iron crank on relationship and the gold crank on discovery information, and that's a very expensive part of what we're doing.

We couldn't possibly defend anything based upon the mailbox domain. It's all hearsay. We couldn't defend it. We can't provide reputation for it, which means it won't stop any of the spam coming in. The PRA bounces around. You don't really know who the mail channel -- what mailbox domain has been checked. You'll

still see phishing. You'll still see spoofing.

Nothing is really going to slow down in that area. We find more people getting more clever on how to gain the system.

I think in terms of providing protection to the system, which is really all it's for, the HELO domain does a much better job of that because you're delegating the responsibility to the MTA. If they can't figure out which customers are screwing up, they don't deserve to be in business, and we're not going to pay attention to their mail, and that's where you have to delegate.

You can't try to decide for the world who can talk. You have to delegate that down to the MTA operator.

MR. BURR: Okay, Steve.

MR. WORONA: I'm Steve Worona, W-O-R-O-N-A, from Edgey Card (phonetic), and Harry, you and I spent a bunch of time on the phone a few weeks ago talking about some issues related to higher ed, and you dealt with some of them up there with forwarding for alumni email addresses, but I actually want to pick up on that, and it's related to the question that came in on the card, and it's further related to a comment that was made earlier this morning to some of the crypto issues and the need for a simple solution for people with small

businesses who are also coming in on home lines.

The issue I want to pick up on is people with multiple email addresses, which I think is more and more all of us, because I suspect all of us at least have a business address and a home address, and if we have an alumni address that goes back to our university, that's three, and if we're hanging on to a bunch of Yahoo! and Hotmail addresses so that we can throw them away when the spammers find them, we've got four or five or six.

So my concern about the Sender ID framework as it now exists focusing on the from address is if we're sitting at home or in a hotel or connected to some ISP somewhere and want to use the single SMTP server that that ISP is offering, which is a well behaved SMTP server which is some sort of read before send authentication so it knows who we are, which I won't say is the dominant approach today, but it's a well functioning mechanism today to allow people with multiple email addresses to send them from a single SMTP server.

It seems to me that the sender IP framework breaks that whereas CSV supports it, and that may not be a reason to go all the way to CSV, but I'm wondering if you can see some way to adopt the advantages of HELO based authentication, HELO based reputation, with

whatever you're doing with SPF to not throw out that baby with the bath water.

MR. BURR: That's to Harry, right?

MR. KATZ: So, Steve, you've isolated a very interesting and I have to say a difficult scenario where people are sending multiple email from multiple addresses through the same SMTP server.

MR. WORONA: Which I think is the norm.

MR. KATZ: I'm sure in some organizations it is, yes, but most -- I think that if I'm logged onto a particular network, I'm sending mail using one particular address most of the time through that network.

Certainly that's the case in corporate environments. It's certainly the case in something like major things like Hotmail. If I'm logged on to Hotmail or I'm sending mail using the Hotmail address. I'm not saying that we can cover all these scenarios in a nice neat fashion, but I think that if a domain is authorized to send mail through -- if you're authorizing, pardon me, a particular SMTP server to send mail on your behalf, the include mechanism does allow you to point to that other domain that is authorized to send on your behalf.

I am not sure if that would cover all the

scenarios, but I think that's what it's intended purpose is.

MR. BURR: Doug, do you want to comment?

MR. OTIS: Well, actually you could build a system that uses a name list of HELO domains to effectively implement the same thing you have now with the SPF record, so if you want to prescribe the mail channel, you would just simply use the name list and that gets rid of having to do with hundreds of DNS lookups. You do one lookup, and you compare the HELO domain and that describes your mail channel, and that allows you to run your PRA algorithm if you would like.

It doesn't stop you from doing what you do now. It would just be a different approach for doing the same thing, but it would also provide a name that would more likely be used for reputation, so that you don't accidentally step on the wrong toes. You don't gore the wrong ox, and that is I think what's really important.

You want to also protect the DNS system. That's very fragile as well. The transaction identity on DNS is only 16 bits, so it's very important to be careful on how you use it as well.

You're dealing with a lot of old protocols that are not robust, so we have to be careful with them.

MR. BURR: Do you want to comment?

MR. CROCKER: I think this last question underscores the challenges in designing anything in this space, and even worse, challenges in evaluating them. There is -- I think it's really easy to miss just how diverse and variable things are.

The amount of computing power, the nature of the access people have, the frequency of access they have, whether it's dial-up or whether it's low speed or high speed, the amount of transaction traffic that can be tolerated or required, the amount of administrative effort, the amount of change in their usage scenarios, whether they're mobile or whether they have multiple addresses and so on and so forth.

The tendency that has dominated much of the efforts to design solutions for the spam problem have tended to identify very popular, very useful scenarios and ignore the rest, and those solutions are useful for those popular scenarios. They tend not to be very popular for other scenarios.

MR. BURR: Right down here.

MR. ANDERSON: Dave Anderson. I think we're missing a very important point, and that is that IP based solutions -- I think we all see enough cases, enough problems that nobody believes that all mail is likely to ever be authenticated using an IP based

solution.

So as a result I think using an IP based solution to exclude mail because it's not authenticated I don't think will ever be a reasonable scenario to go through. Once you've said that, hey, this is still a good thing, there's a whole bunch of mail, as you said, David, that many very popular cases where this gives me a tool that I can assure my mail is in fact authenticated.

I can use a VPN to go back to my home site and thus send from my home site so that things work out just fine, and in answer to your question, is this really just an expedient, well to some extent it is. We need to get something out there. Frankly I think broad presence of an IP based solution is more likely to attract people to do a more idealized signing solution than any other single thing we could do. We need to get moving and then this will take us to a more ideal case.

As far as long-term, long-term I don't think you can dismiss the forensic positive effect of having two different ways to authenticate something. Yeah, I may not be able to cover every case of an IP based solution, but if I have two ways to do something and one of them is not working, that's usually going to be a good place to go look, and especially if the IP based solution does

work and the signing solution does not work, that's going to give us a real clue as to how to go fix the highly variable environment.

So I think you're looking for some redundancy. There are two cases that I think can cover a large number of the cases we see out there. We're not going to get perfect coverage but I think we can get very rapid adoption. Thank you.

MR. CROCKER: You're looking at me. Boy, I'm speechless. That's really tempting to say, but not this time. There's a peculiarity about CSV that's really easy to miss. It's usually included in the IP schemes because it uses IP and the current version of the spec to do the authentication part, but it's actually doing also authorization and accreditation, and it uses domain names for that.

In fact the first version of CSV that I wrote allowed multiple forms of authentication, and we took that out because it was really confusing people. The reality today for immediate utility is that IP based validation, IP based authentication is just very convenient.

On the average people don't think it's a very good, long term approach and I concur with that, and so our view is that as other schemes are convenient to use

for authenticating the domain name that's used in CSV, that can be spliced in really simply. I don't know how easy or difficult it is to splice it into some of the other schemes.

MR. OTIS: Can I add to that? Right now we have a model that's working. We have essentially an IP based reputation system that's widely deployed. It's widely used and it's fairly effective at protecting the network resources heading into the mail system. It's not perfect. It doesn't get rid of everything, but it gets rid of quite a bit.

And I think that role is going to be needed in the report long into the future, especially if you're looking at more intense ways of ensuring the actual originator where you're using signatures, that resource is not going to be protected by these schemes, so you need effectively two levels of protection.

I think analogy would be the garden gate protecting the path to the front door. You still need the front door, but you also need the garden gate, so we have a model that works, and that's based on IP, and I'm saying that as we move into the name based reputation services, we need a reasonably strong name that we can start using to get a reputation database ready for the front door.

So I think the only strong name that we have in the mail channel unfortunately is the HELO domain and it needs to be fixed. When we fix that, then we have a directly verifiable name that we can use to start building on that database. It starts at the front gate. Now, we have to verify it. We don't trust it that much, but now that we have that database we can use it at the front door.

Unfortunately I don't think you can use any of the information you're getting back from Sender ID or SPF for that because you simply can't trust it.

MR. BURR: We'll take a question here.

MR. BARCLAY: Hi, Doug. This is more a clarification of your statement that HELO is the only domain you could build a reputation on. I'm sorry, Robert Barclay, B-A-R-C-L-A-Y.

A relatively common case that at least I've observed in my independent email, and I'm sure other people have seen in the real world, is that what I will call moderately bad or not quite completely evil spammers will send using their own domain but through a variety of network providers until they either get reigned in or kicked off of each one.

If the domain is only based on the -- if the reputation is only based on the HELO domain, then each

of those network providers will be damaged by that sender, but doesn't that bad sender deserve their own -- is it your assertion that we don't have a good way to give them a reputation or that we shouldn't?

MR. OTIS: No, as I said in the mail broker or the analogy I used was in the insurance industry, the broker is going to be responsible for knowing who the good mailbox domains are. In other words, that's their job, and they're going to have to do a clearing house. They're going to have to figure out a way of working among themselves like the insurance companies do to know who the bad actors are and to keep them from getting the customers.

It's their job to make sure they get rid of their bad customer. If we somehow magically implemented Sender ID with perfect security and we established a reputation system on it, what would happen is they would all move into the large domains. We would be left with the same situation.

So you still need to weed them out, and the only people that can weed them out is the MTA or the domain operators, the mail systems that allowed them in. There's where you close the door.

MR. BARCLAY: Doesn't deciding to allow them in imply that there's already a reputation system to make

that decision on?

MR. OTIS: The reputation is going to be on the broker. You can't base the reputation on hearsay. You can't trust an unidentified broker that someone may or may not have authorized, right? We don't even know if you've been authorized for a particular field because you don't even know what fields they were trying to authorize by the records.

It's a very messy situation, so you're basing it on hearsay. You don't know if the MTA has been compromised. You don't know the different administrative regions it's gone through. You don't know who may have gotten the information as it headed towards you. There's nothing that you can trust, but you can trust that you know the machine that's sending mail to you, and because you know that, you can base a reputation on that fairly verifiable information.

Everything else is just too flimsy to trust a major lawsuit in terms of staking your company's future on saying, yeah, they're bad. Well, I think they're bad. Maybe they're bad. You can't do that.

MR. CROCKER: There are a lot more author domains than there are MTA domains, so there's a degree of scaling benefit that you can get from something like HELO validations, in addition to which there are

aggregate problems when we have the bought networks, with 60 million machines compromised, the individual operators of those machines, the owners of those machines probably are not real good at fixing things and probably don't even know there's a problem.

The aggregate performance of a bought network on an operator's network is probably visible to that operator if only they are given some feedback. CSV provides a way in which the aggregate reputation of those misbehaving machines funnels into the reputation of the operator of that.

One last point, HELO can vary. You can have different domains put forward according to different senders, if the operator chooses to behave that way.

MR. BURR: Harry, were you trying to get a word in here?

MR. KATZ: I wanted to make the point that in the case particularly of the bought networks, I think you have mail emanating from these networks purporting to be sent from a huge variety of domains, all coming over the same set of machines, and if those machines have, for example, published CSV records, they will appear to be perfectly fine, and yet all this spam and spoofed mail will have been emanating from them.

So I think I would reiterate my earlier point,

that it's fine and dandy to go and authenticate the specific machine that is sending mail. I just don't think it takes you far enough. I don't think it's frankly accurate to suggest that this is -- that the Sender ID identity that we check is hearsay or untrustworthy whereas the HELO domain for some other reason is.

I think they're roughly comparable in their degree of reliability, and I don't believe fundamentally that we can simply dismiss this just like I said it doesn't take you far enough.

MR. CROCKER: I agree with you, Harry.

MR. BURR: All the way in the back there. We'll get around the room here.

MS. OLSON: Margaret Olson. I guess the question I would have for Doug is that although I completely agree that there is value to holding the channel accountable, when you talk about the channel essentially -- the MTA operator enforcing, knowing who their customers are, knowing if they're good or bad, what you're essentially saying as far as I can tell, correct me if I'm wrong, is that everyone that operates an MTA needs to know trade information about customers so that if someone got kicked off of service X and they come over to service Y, the service Y has no way of

knowing unless there's some kind of clearinghouse that rates people might like a credit rating.

I guess I find the PRA approach to be far more transparent to the sender and a far more gradual way to accomplish that, because ultimately I think what everybody here today has agreed on is that we need to hold senders accountable, and authentication is the first step to doing that, but you have to know who that sender is, right?

You can't just say to the people operating the mail servers, Guess or call up every other one and ask them if they kicked these people off, right, and that's what to me is attractive about the sender based -- the PRA and the IIM and DomainKeys because they concentrate on the people who are actually composing that now, and gives them a reputation. Those are the people who -- that's where the reputation needs to be.

MR. BURR: So let's have two quick answers, and then it will be time to call it.

MR. OTIS: In terms of scales of problems, the number of bad actors really isn't that many, so in terms of scaling out knowing who the real bad players are, it is not a long list, so I think the players in the industry can figure that out.

The people that don't know that list, don't know

who the bad actors are, they're only recourse is diligence, and most of the serious mail providers carefully monitor their SMTP air log and notice the bad actors and move them off the system.

They learn that way or through a type of clearinghouse or industry scuttlebutt or however you want to describe it, but they know that they're not going to provide access to a certain group of people or they'll monitor the system and see when that happens and kick them off.

It's their responsibility to run a tight ship, and we can't say we're going to trust anyone and everyone that sends mail that has been authorized, may be authorized or we think they're authorized and say that, now we're going to give them a reputation because now you're not allowing the person that owns the mailbox domain to defend it because you haven't given them any method of defending their mailbox domain which is very important to them. It's how do you defend that?

MR. BURR: Harry, is there a counter answer succinct here?

MR. KATZ: First of all, I would say if the number of bad actors was so small and they were so easy to find, we would have knocked them off already, and I think the evidence is that if they are small, they're

extremely crafty and move around and change domains and IPs and networks all the time so we do need I think some solutions to attract them wherever they are and under whatever domain name they're sending mail.

Another point that is sort of the converse of this is that we want a mechanism that allows legitimate senders ways to protect their domains from spoofing, ways that they can distinguish themselves from spammers, ways that they can demonstrate their bona fideness, and we think Sender ID allows them to do that by allowing them to publish records that clearly identify themselves as the domains that are sending these messages and are identified in those message as being legitimate senders.

MR. BURR: Thank you all, panelists, for your time, and I believe we've due back at 3:15.

(Applause.)

(Break in the proceedings.)

PANEL 4: EMAIL AUTHENTICATION METHODS:
TESTING, IMPLEMENTATION AND EVALUATION
MODERATOR: SANA D. COLEMAN, FTC
PANEL MEMBERS:
SCOTT BROWN, ColdSpark
MIKE CHADWICK, Go Daddy Software, Inc.
DAVID FOWLER, @Once
CARL HUTZLER, America Online
KARL JACOB, Cloudmark
BILL KARPOVICH, Port25 Solutions
BARRY LEIBA, IBM Thomas J. Watson Research Center
DAN NADIR, FrontBridge Technologies
ROBERT SANDERS, Earthlink
RON SCHNELL, Equifax Marketing Services
RAND WACKER, Sendmail, Inc.

MS. COLEMAN: Thanks for settling in. This is the last panel of day one.

(Applause.)

MS. COLEMAN: What we've attempted to do with this agenda is to really take you through a journey of email authentication. We started out with the basics, what does it mean, why is it relevant. Then we took you to a discussion about the policy framework, in other words, what were the boundaries that we need to look at

when examining this issue, and then we gave you presentations about domain level email authentication proposals.

So this panel is going to talk about, where we are with these proposals. Have we tested them? How have we tested them, and what have those results shown us? So this is going to be very exciting, and the panelists have promised me that they are going to be as entertaining as possible, so sit tight.

They're going to come up one by one, and if I may just go ahead and read the names of our distinguished panelists: We have Scott Brown, CTO of ColdSpark; Mike Chadwick, Vice President, Application Development of Go Daddy Software; David Fowler, Director of Deliverability and ISP Relations @Once; Carl Hutzler, Director of Anti-Spam Operations, America Online and he brought his fan club, okay, nothing wrong with that; Karl Jacobs, CEO and Cofounder Cloudmark; Bill Karpovich, SVP Marketing and Strategy of Port25 Solutions, Inc.; Barry Leiba, Senior Software Engineer, IBM Thomas J. Watson Research Center; Dan Nadir, Vice President, Product Management of FrontBridge Technologies; Robert Sanders, Chief Architect of EarthLink; Ron Schnell, Vice President, Equifax Marketing Services; and last but not least Rand Wacker,

Director of Product Strategy and Planning, Sendmail, Inc.

(Applause.)

MS. COLEMAN: Scott, why don't you come on board and get us started here with your presentation.

MR. BROWN: Being a Brown, I've always been first. We'll start with the thumb trick, right, everybody knows that. I'm trying to keep it active. All right.

We've heard a lot of the background information on SPF and Sender ID and all this stuff so I'm going to fly through a lot of this. I just wanted to say that from ColdSpark's perspective, everything kind of happens at the margins, so if we can get a 3 percent or 4 percent, 5 percent switch in the spam or the fraud that's out there, we're doing pretty well, and I figure being in Washington, D.C., a 3 or 5 percent switch makes sense. I am trying, guys. Work with me here.

So at ColdSpark what we looked at is really kind of thinking about the SPF, Sender ID versus the cryptographic. We are a big fan of the cryptographic solutions. We do a lot of work in the financial space, and so being able to actually sign a message and provide a measure of validation in the transport of the actual message is important to us, also being able to drill

down right to the user level.

So a lot of what I'm going to talk about today is focused really on the results of our testing with the different cryptographic solutions that are out there.

So when we looked at our implementations for both DomainKeys and IIM, we had clear challenges right off. My lead engineer says, "if you're doing crypto, it's a CPU problem," which everybody has talked about today.

Obviously there are DNS hits whenever we do these lookups trying to figure out all the records, so we had a couple of clear goals. One, that it had to have a low overhead. It had to be fast. We had to have a really high speed DNS resolver because of the potential DNS lookups, and then going back to what everybody has talked about, there has to be configurable outcomes so if it fails, what do you do with it.

Do you block it; do you accept it anyway, flag it; do you throttle it down, apply some type of quality of service to the actual transport layer, things of that nature, so with those things in mind, we actually put it together and ran some tests, so I want to talk about what we consider our real world tests.

In our lab we have a set up that does actual full DNS lookups that delivers to mailboxes and mail servers that look like the real world. Some are slow.

Some dropped DNSs. Some do retries. Some block all together, so it's really trying to mimic the Internet in our little lab.

We ran this test on your basic $2,500 Winnex box, dual xeon, on two giga RAMs, like I said, lots of domains, full DNS lookup, and this is a JAVA based solution so some of these CPU numbers are going to look high because it's JAVA based. Welcome to my world.

So the baseline right across the top, you'll see that our base Spark Engine running real world is going to do about a million messages per hour, inbound and outbound, with about a 30 percent CPU hit.

When we add-on IIM, our CPU went up pretty high, and we attribute that to the JAVA based implementation. However, it didn't change really the speed at which we were able to transmit messages. We were still able to get well over 800,000 messages per hour going through our server using that crypto.

With DomainKeys, it was actually a little bit faster because we only had one hash that we had to run. The IIM actually had a double hash that we had to run, and that gave us a bit of a hit in JAVA, so that DomainKeys ran slightly faster.

What's interesting is what happens when you put this into the lab without the real world scenario. So

when we do a straight high capacity, smart host throughput so that we're not doing all of this slow downs and bounces and just pumping messages straight through, what we found is that we didn't gain much in our implementation, again, same implementation of the technology.

It topped out around 850, 950, a thousand messages per hour. That's still way beyond what most people are trying to do on a single server outbound, so in our estimation, we feel like both of these solutions are effective and can work for a corporate environment, and really that's kind of the key that we're looking at here.

By pushing it under significant load, we found that we can get this kind of speed, 800, 900,000 an hour and still be able to run efficiently.

So my outcome is pretty easy. We think it's practicable and effective. We like the crypto better than the SPF type or the path based. We think that the performance impact can be minimized, and that if you can actually run 800 or 900,000 messages per server per hour, outbound or inbound, that's going to certainly cover what people are capable of sending or require from a single server.

And then again adoption/roll-out, being able to

have those configurable outcomes so that you can block, accept, flag or slow it down. That's kind of what we're thinking about.

Thanks.

MS. COLEMAN: Thank you.

(Applause.)

MS. COLEMAN: Thank you, Scott. Next we'll here from Mike Chadwick.

MR. CHADWICK: You all know who I am now. I work for Go Daddy.com. I'm going to skip a couple of these early slides. Go Daddy is a small company. One of the unique things about it is that we serve well over 2 million small businesses, and our email system is fairly large where we have well over 3,000 domains that we manage, and that creates a unique set of problems for us in this industry versus someone that's working at corporate, large enterprise consumers or companies.

We have a different set of issues we've got to do, so we really looked at our implementation being very multi tiered. We already have in place all of our own spam filters we wrote. We subscribe to the Bonded Sender whitelist. We have our own large blacklist that we run, and that stuff is not going to go away. No matter what solution we adopt here authentication-wise, we can't let every cache come into our system.

There is no way, we would have to create the quadruple or hardware or more than that. We handle -- we block probably about 60 to 70 percent of all connections coming in today at the IP level, 60 to 70 percent, a very large percentage.

We cannot just open that up and say, "Okay, now we're make going to check emails coming in to authenticate them." There's no way. We currently support SPF Classic. We rolled that out a few months ago, and I'm going to go through some stats we have related to that a little bit later.

We chose SPF for a couple of reasons over crypto. For us to roll out the crypto solutions, we have to basically put in a private key management system for 400,000 plus customers that are going to use our email system to send email, and that right now, I didn't want to do it this year so we're at doing it sometime in the future.

There's a whole set of issues around that because people are giving us their price, and secure those, how secure do they have to be? Do we have to get HSM for them and that sort of stuff. It's a much more complex issue for us than just rolling out SPF and relying on our customers as you usually publish your own SPF record using some of our tools.

Obviously we want to keep things here for authentication. We believe everybody has a right to be able to protect their domain, no matter how small. If you're a small business, you're running a flower shop, you have two employees, you have a right to protect your domain and be able to prove that you are the right person to be sending from this domain because a lot of times you'll find -- I have friends who have small businesses and that they get thousands of bounce backs a day from people just using their return address to send out spam all the time.

And that's the problem we definitely want to see fixed as soon as possible to help prevent our customers that deal with all those kind of bounce backs and spam they get that's just really out of control right now.

Some of the hurdles that we have come into, for us we're kind of key where with SPF right now, it's been out there now for quite a few months, there's no real centralized testing process no validation testing process. Large corporations have been -- we've been blocking their email or rejecting the basic SPF that they misfigured. We get on the phone with them. We walk them through it. We change the configuration. There's no real process for rolling this out that's clean.

Another big issue for us is we forward literally millions of emails a day. We're just a go between. They'll apply for a domain with us. They'll want it to go to their home account or whatever it is. That stuff gets forwarded to us. We do millions of those a day, and the current petition doesn't support that very well. It puts a lot of burden on us to do some additional checking, whether we do it in spam filtering, virus testing, whatever it is which increases the load on our systems.

So for us, ideally, this is in the ideal world, we would choose one solution for the next year and a half to two years, whatever it is, that's what we roll out. If the industry adopts three or four solutions, our customers are going to call us and say, "We want that one, we want this one," so we'll be forced to have every single one of those, and our system gets much more complicated.

It's important, Jason over here, my lead engineer on this system, he has to go out and do things with his team, and it just gets more and more complex, creates more issues in production and we're going to bounce more through emails in time. It's just going to create more issues, so for us ideally start with the simple approach, pick one that we all agree on as the

best approach to start with and roll it out, see what happens for a year or two, see how it works, see how well spammers get around it and then kind of tweak it out from there and then roll out other solutions as they're needed but not trying to solve every problem with three or four solutions at one time.

Obviously we're committed to supporting any approach. We're going to have to. Our customers will make us, and we're also very committed to Sender ID. We Rolled out SPF today. As Sender ID application moves forward, we're going to support that. For us it's a much easier solution. It solves I believe 90 percent or so of the issues out there so they're really helpful.

Some the small staff starts. Like I said, we currently block about 70 percent of all connections coming into our system. Our implementation right now, SPF, about 7 percent of all email coming into already has published SPF records. Basically 18 percent of email checked against SPF records. Email is coming in either from a spammer or somewhere else and we're actually rejecting those emails, and we're doing what they tell us to do, okay, reject it, and we reject a lot of emails that way.

About 14 percent domains that pass our checks are actually known spammers listed on some spam list

somewhere, and that's actually increasing, and we don't really know how many of these emails were actually passing SPF or anything else that are actually spam. We don't have good numbers for that right now.

What it basically shows though is that spammers have no problem finding a domain, publishing the records and getting spam because it's really pretty trivial by domain.

Back to my last point which I've made many times before in the past, is that these systems are pretty much useless without some kind of reputation and reputation really has to be controlled that come to the point of purchase or transfer of ownership domain.

Otherwise, it's just going to be something pretty easily abused by spammers as they get into the kind of reused domain market. They watch what's going to coming through. They buy it that day. They start spamming that day. It still has that domain that has a very positive reputation associated with it so it's key that registrars get more involved in the reputation process to ensure actually that there is valid reputation out there, and it's delayed, and we also forward people that are buying domains that give us good information which will help all this stuff.

(Applause.)

MS. COLEMAN: Thank you, Mike, and now we have David Fowler, @Once.

MR. FOWLER: So I'm the first email services marketing person up for the day, so hopefully you won't be asleep or I won't be directing myself or taken myself out of the missile path as they come over here.

So my disclaimer on the presentation is I have my daughter doing a quick spell check on that so if you see typos, I'll certainly make sure she hears about it later on this evening.

Really quick, sort of moving forward, I had timed this about for about an hour and 20 minutes but I certainly want to give everyone else on the panel the ability to come up here, so I'm really happy to be at the Federal Trade Commission.

My name is David Fowler. I work for a company called @Once, a corporation based out of the Portland, Oregon, as you can tell, and we'll talk about @Once corporate environment. There will be no shameless self-pitches here today, so put your seat belts on.

The evolution of email marketing is really an important key element because it's really our livelihood, right, and I think from just a marketing perspective, I'm going to show you some of the things

that you've seen around authentication.

We also are IP and SPF compliant as all our clients are as well. I'll talk a little bit about the business challenges and the compliance hurdles and the @Once efforts for authentication adoption.

Again we're based in Portland, Oregon, founded in 1998, 60 employees and 40 clients, and a drum roll please, we're actually profitable which is good news.

We do everything email and everything around email, so if you subscribe, for example, to some of our clients who include Nintendo, Niki, Warner Brothers, Home Shopping Network, Cingular Wireless, those types of email communications are coming out of our shop based on the tenth floor of the 900 building.

Here we go again. Email has evolved from technical placing, but more importantly, the value being delivered to the consumer with more relevant and more personalized messages has evolved over the years. I don't think any of us would disagree with that.

As email has evolved, companies have seen more value and return being driven from it so that the challenge becomes the critical component of driving revenue for companies. In some cases almost 30 to 40 percent of a company's revenue comes from permission based CAN-SPAM compliant, email marketing, and the last

time I checked we weren't breaking the law for doing that, so that's good news.

With the complexity of consumer value and company value rising, the company's reliance on the challenge has grown exigently so that when basic things, like, can I deliver emails to my consumer who requested it comes into question, it's a big deal for clients out there.

You should not be able to state that for a large company email marketing is a critical channel for business success, and while it may not be a big issue for my parents and myself to have one email be mistakenly blocked, it's a huge deal for a company that has their revenue consumer life cycle value tied to that mechanism.

We've been following the Email Authentication ups and downs over the last year very closely, and I believe it's time for widespread adoption, get on the playing field, put the kids on and start the game and hopefully we've done that.

@Once is SPF compliant. With our technology platform, I find it rather simple actually with no significant major business hurdles to overcome. I think the biggest challenge we had was to decide what flavor pizza and beer was going to be delivered to the

technology guys and gals that actually do the coding itself.

So for us we obviously have a lot of resources available to us, which may have not be the case for a small or medium sized business so that potentially creates some challenges in that realm.

With that said email authentication solutions can pose several challenges to those who do not have the necessary and general resources who are not fully versed in the technology requirements.

Permission based email is still about accountability, and authentication still does not guarantee delivery of email through recipient's email. There are still many other factors that have affects on that issue.

I don't have much light so I apologize for that. Correct two way communications still remains a challenge to the senders and receivers of email.

There are numerous policies, both internal and external that an ISP can implement to control the flow of email into the networks and quite rightly so, so from our perspective or ESP's perspective, it's a case of the old Ghostbusters and with my best American accent, "Who are you gonna call?" All right. Not enough caffeine in the room.

Okay. With no consistency, that leaves the guilty until proven innocent approach, while valuable to the spammers, does not create a level playing field for the legitimate senders of commercial email. We still have a long way to go to erode the one-sided accountability playing field.

Email authentication is a major milestone in addressing the spam problem. It will not solve the spam issue, but along with legislation and industry forming good, best practices, it's a necessary and valued first step.

The challenge remains that in order for businesses to adopt rapid authentication solutions there needs to be a consistent standard and support for these solutions from the ISPs and business community. We have to work together. We can't be on different teams, ladies and gentlemen.

Resources should be made available to businesses that adopt authentication and aggressive public awareness campaigns should explain in detail the issues surrounding authentication and the expectation for email delivery.

My expectation today is if I stick a stamp on an envelope, it gets to where it's going to go, and the same should be applied to the email world.

So @Once has demonstrated that we've completed early adopted authentication solutions and will continue to support the cause, working directly with our industry association buyers of the like ESPC and a few others involved, we will continue to educate our clients and conduct the appropriate and necessary training to support email best practices.

Thank you for your time today, and I look forward to your questions.

(Applause.)

MS. COLEMAN: Thanks, David. Now we have Carl Hutzler from AOL who is going to give an overview as well.

MR. HUTZLER: Good afternoon, everyone. I'm going to give you a quick overview of what AOL is planning to do in the authentication realm, and specifically what we plan on testing, because we really don't -- we don't have a technology. We really don't know which one is the best. We're sort of looking at all these as addressing a sort of different tact on each of the authentication and verification areas that we think are needed.

So we plan to test many different types of authentication technologies, and I'll take you through a couple slides that show you which ones we have immediate

plans for and which ones we are looking to do early next year. We think that testing is critical. We're scared about the Internet mail backbone. I'm more scared sitting through some panels today, especially the gentleman down there that has five email accounts and is sending out through Comcast.

I do the same thing myself, and I know I have to change that practice, or maybe I don't. I don't know. We'll have to see which one of these applications ends up being a winner.

Testing will identify a lot of situations we think where these proposed technologies may break the existing infrastructure, and more importantly, the things that they do break, how big are those things? Are we talking about 99.9 percent works just fine and we have a tenth of a percent out there and there's one MML marketing thing that needs to change, or are we really talking about 80 percent works and there's a huge gap of mail that doesn't meet these criteria.

We're going to be implementing these things in what we call a dry mode at AOL. We're not going to be affecting mail with them. There's a chance we might. If Citibank calls us and says, "We are getting hammered by phishing, we want you to reject everything that's not SPF compliant for Citibank," we may do that, and we'll

caution them that forwarding and other things where SRS isn't implemented or PRA isn't implemented might break, but I think 99 percent of the time we're not going to be affecting mail so don't panic.

We're going to try to look at -- we are going to look at all the metrics that we're going to get out of this dry mode. How many domains are publishing SPF, how much mail does that represent, how much checks out, how much doesn't check out, what are the situations where it doesn't, and we're going to be doing that as you'll see for a lot of different technologies here. What operational issues are we going to encounter?

I think you heard a little bit from Go Daddy's software. They have all these domains they have to work and what a pain that is. Thankfully I have a lot of mail but only three or four domains I have to worry with.

There are other operational issues. We've already found -- some of the folks in the room may have remembered, I was saying we would be probably be doing SPF and Sender ID inbound checking in the fall. We've actually found a couple of implementation issues in our own software development trying to implement these technologies.

Not that the technologies themselves are broken,

but just developing that for our own infrastructure, we found a few things that didn't scale for our platform, and a few DNS caching things we had to work through, so we've had a little bit of a delay in doing that but we're getting close.

Also obviously suggesting areas for improvement to these technologies if we're smart enough to recognize what those are. I don't think we probably are. I think the guys in the room are probably smart enough for that.

So here's our test plan. Part 1, these are the I guess IP approaches or path based approaches, if you will. The SPF Classic, we've actually been using for awhile now, since July. Brian Barrious is in the room. He actually implemented a form of automatic whitelist updating for certain well trusted domains that AOL maintains a whitelist for.

We're actually using SPF records so that those domains that we trust can update their own records, and we can feed that in as opposed to constantly having to work with Mark and Miles to know which new Yahoo! group servers were added and things like that, so we started doing that.

That's certainly a use of the technique I think very few people are thinking about, but we saw it as valuable to us.

In late 2004 or early 2005, we hope to be testing all of our inbound mail in a dry mode again, for this particular SPF check. We will not have SRS checking enabled in that first incarnation.

Sender ID framework, you've heard a lot about this in the news. We are now publishing SPF, not only version one record, the classic, but also version two. We're also going to begin checking the 822 from domain. We're not going to be checking the PRA algorithm initially. We're just going to be checking the domain against the SPF V.1, V.2 records.

It's only a partial test, but we think because there's not a whole lot of domains signing or using the PRA on their outbound systems it's probably a reasonable test to do at this point. If we start to see that adoption rate go up, I think we're going to have to switch over and start giving PRA as well.

Part 2 of the test plan is looking at the signing based approaches. I probably should have put CSV on the other page and BATV on this page. I apologize for that.

DomainKeys and/or or Cisco IIM, we're looking for ways to implement outbound signing on our system. We thought initially we might be able to do it because we do use Sendmail, on our last sort of hop getting

out of AOL, but we found talking to our architects that the way we use it is pretty strange, and we're not able to just sort of use the implementation the reference implementation that's been put, and for $14 an hour, no, for 140 an hour.

So we're looking at that, and we're hoping that we can sign outbound mail very early in 2005. The folks at Cisco just came up to me today and are interested in trying to get us to do it on our outbound system. We're probably going to be working with both organizations to see how we can do it. If we can do both types of signing, we would like to do that as well.

Client SMTP Validation, again I probably should have put this on the first page, because it really isn't a signing technique. We're going to be implementing this along with SPF and Sender ID checks although in a little bit of a modified approach. We're going to use the SPF 1 record to compare the HELO domain. It's not exactly the way the CSV implementers had envisioned this, but it should be an interesting check to tell us how many people might adhere to this just using their current HELO.

I know AOL, when we send outbound mail, for AOL.com, we HELO as AOL.com. There are probably a lot of domains that naturally fit into that in a very

simple case.

Until we start seeing CSV adopted with the new record type, we don't really see a need right now to start looking at that on our inbound side, so again it's kind of the cart before the horse, chicken and egg type thing, and we'll probably look to implement that new record type as soon as we start seeing people adopt it.

We also may try and compare the CSV records and those domains to our internal reputation systems. Everyone here knows about Scomp, and if we can start comparing things instead of just by IP but to actual domain, I think that would be some very interesting data that we could share with the technical community.

So which technology will win? There are a couple of people here that have placed their bets on different things. We really don't know. If we had to bet, we would probably bet on the safe side of things, probably all the technologies are going to win in some shape or form.

Will we ever be able to reject mail that doesn't pass Sender ID? We don't know. Maybe not. Maybe so. Will we see DomainKeys or IIM adopted on a wide scale basis across the entire Internet? You might see it at some of the bigger domains.

I'm not sure how long it's going to take to get down to a small ISP in India, for instance. So we're sort of putting our chips down on betting all across the board hoping that we can implement many of these things, and I think as a big ISP, as a big receiver of email we owe it to the community to do that, and we'll probably have to implement all these technologies in one shape or form.

Testing is critical. Anybody that thinks they can implement these things at large ISPs like an AOL or Yahoo!, et cetera, and start rejecting mail based on this is -- I won't even go there.

While it's impossible to predict the future, we're hoping that the test results that we can provide back to the community will help people who are designing these technologies and implementing them understand a little bit better from a big ISP's perspective what we're seeing out there. Thank you.

(Applause.)

MR. HUTZLER: I have one more shameless plug, if you have a blocking issue or you want to contact me, there's my stuff. If you want to see if your own network is a source of spam, sign up for a feedback loop, and when you do get blocked, if you do get blocked, and you want help, give us a call. There's the

phone number for you. Thanks.

MS. COLEMAN: From one Carl to the next, so we have Karl Jacobs.

MR. JACOBS: My name is Karl Jacobs, and I have two pieces of good news. You're about halfway through this, so we're almost on the other side of it, and we have a completely different way of thinking about this problem because our job is to protect you all from all the terrible things you've been hearing about today, fraud, viruses, spam and all those bad things.

I'm going to talk a little bit about our product set and how we're integrating these kind of authentication technologies into our product set because I think one of the important pieces of adoption here is that people's networks who we are protecting adopt these technologies and we adopt these technologies as well.

So talk a little bit about safety bar. Over a million registered users. Why is that interesting? Well, because it's a peer to peer network that relies on two things, trust and reputation to determine what is and what isn't spam. That will become very relevant when we start talking about reputation around Sender ID and authentication mechanisms.

Exchange server which is designed or Cloudmark exchange edition which is designed for small

businesses. Cloudmark rating which is a content based reputation system, so Cloudmark rating it's underlying technology has been around since about 1998. It processes about 430 million messages a day and about 15 reports a second.

So as far as people who are getting reputation data about what's really going on out there, we're seeing quite a bit about it, and a little bit about what we're doing at the Gateway because there's radically different problems and issues from implementing these problems at the desktop versus the gateway.

So safety bar is an Outlook, an Outlook Express and Lotus add-in technology. The first question, and this has been raised in some of the other panels is UI issues. From our perspective the reputation in our network comes from people voting on the content.

From the reputations that are being done around Sender ID and other authentication mechanisms, the reputation comes at a wider level, and here's kind of the corollary I have or metaphor. If you think about Sender ID and SPF as ways to authenticate domains, one way you can contextualize that is to think about your mileage plan we all have: United Airlines, American Airlines. I like and trust United Airlines, so when they send me a piece of mail, they also send me a

whole bunch of stuff I don't want.

So the UI issue here leave what do we deliver to the user and what choices do we give them as far as things they can block or not block. I don't necessarily want all of Amazon's marketing email about the book club and the movie club and all that, but I do want to get my statements about my account or I might want to know about my Amazon order.

There are spoofing issues. A lot of kinds of conversations happened around this, but our belief is that in all of our products we're going to have to attach reputation to all of the authentication data that we get. There's not going to be anyway for us to make a determination about a particular message without reputation data.

All authenticated email or give users the choice? Part of this is understanding that at the end of the day the final arbiter of what they want or don't want has to be the consumer. It has to be the person receiving that mail.

The idea that we can arbitrarily decide further up the stream what they should and shouldn't get is a little much. Can we get rid of a lot of the bad stuff and authenticate a lot of the good stuff? Yes, but at the end of the day I think we need to think about the

consumer.

So how does this look in a user interface? I hope you can see all this. If you look at the upper left-hand side, you'll see a block, spam, fraud button, that is our feedback loop into our system so we have millions of users out there hitting those buttons every day.

If you look further down, there's my rating which is the reputation for the person submitting content, meaning do we trust you or not submitting content into our network, and then you see a little thing called Cloudmark rated, so Cloudmark rated is the rating system that I'm talking about, and in fact it's using a couple of things to make the determination in this case.

It's using our reputation system underneath and the content based reputation, meaning on a per email basis. That means that I could say, I want Amazon's book list and I don't want their movie list. It's also using Sender ID and other authentication mechanisms at a higher level to understand what the gross level of input in the system is, meaning is this somebody I should trust overall.

And lastly we're using a lot of that information to give something to the user so they can make a more

informed decision. One of the big issues here and it actually hasn't been discussed is that a lot of consumers don't understand what's going on in the systems. They don't understand why something is being blocked. In many cases they don't even remember signing up for these things, and so communicating that to the user is going to be critical.

So now we're going to shift gears a little bit and talk about integrating these authentication systems into the Cloudmark rating. As I mentioned, it's a reputation system for legitimate senders of email.

One of the unique characteristics of this is basically that it's a feedback loop. Not only do we broadcast the Cloudmark rating to anyone that wants it, but if you're a sender of email, you can actually go to our web site, look yourself up and see what emails have been blocked or not been blocked so that's a critical piece of the feedback loop that people need.

It's been extended to support SPF and Sender ID. Right now you can come to our web site and you can download an SDK that allows you to do a check against reputation as well as a check against SPF, et cetera, so basically you look up the authenticated domain and then you can look up the reputation.

In our mind this is the key critical factor in

making sure that these are successful. The reason being we have plenty of authentication mechanisms on the web, in email and in the real world. The problem is they don't work very well unless you establish some type of reputation around them because you don't know who to trust.

We leverage the same DNS based architecture of SPF and Sender ID so the information can be gotten in the same way. As we mentioned we're going to check authentication and reputation. One of the things we're doing in our reputation system is trying to provide additional data so you get a rating that is essentially zero so a hundred percent, the people who think this is good, a confidence, meaning how confident we are and their status in the system.

There's a whole bunch of other pieces of data under that. One of the more interesting ones is velocity, so where is their reputation trending over time and how quickly? Are they rapidly decreasing in reputation which is probably someone you want to hold up or are they rapidly increasing in reputation which means you probably made a mistake and a bunch of other people are voting in the other direction.

So the last is our Gateway products. At the Gateway there's a whole new set of challenges for

dealing with this. One, do you drop the messages or tag them? There's been a lot of talk about, well, if they're authenticated, then they're probably good. We heard that's not the case. Spammers use these things as well.

Probably best to tag them at least initially as I think a lot of people are doing to communicate the information to the end users and to the administrators but not do anything with the messages itself.

The biggest question we are asked I think as a company designed to protect consumers and enterprises against spam is, should we override the spammer fraud decision, meaning if I'm on the Sender ID list and I'm authenticated, will you override all your controls and let me through, and the answer is absolutely not.

There's just no way this early on that we can trust that those systems were going to be secure against a lot of the attacks that we see. Reputation systems will help a lot. The jury is still out as far as opening up our networks to that kind of inbound messaging.

The last thing I want to talk about is again this topic of integration with per user preferences. The idea I think that again at the glittery or anywhere upstream we're going to decide what consumers should and

shouldn't get is going to be problematic, so it's really kind of a battle between what the user wants, what the corporate policy is at the company or the enterprise and what the sender wants to accomplish.

And again we think a lot of the solutions in this space are going to be around feedback loops that allow senders to do a better job and see what's happening. They allow corporate policy to be set that consumers can understand, and at the end of the day, if the user wants it, they allow users to set their own policies about the kinds of things they want to see and they don't want to see.

So we think obviously authentication is a value part of overall email defense. Reputation we think is the key piece. Authentication is something that we would like to happen very much because we think reputation is going to make a big difference in this war against the spammers and fraudsters.

In our minds protecting employees and consumer rights is a must, and this kind of goes to the argument about kind of the little guy versus the big guy.

In many ways, the more we work on systems that solve the larger problems, the harder it is to satisfy everyone, and while we actually think that we'll have a positive overall effect on email as a medium, we have to

be careful not to take away all the reasons that we use email in the first place.

We're in the middle of real world testing and deployments underway. We don't have a lot of the great data that everybody else has because as we're integrating these into our larger customer's networks, making decisions on these types of things is a lot more scary for us than others who are just out there trying to collect the data. That's it. Thanks.

MS. COLEMAN: That was Karl Jacobs. Next we're going to hear from Bill Karpovich of Port25 Solutions.

MS. KARPOVICH: Good afternoon. My name is Bill Karpovich, and I'm SVP Strategy and Marketing of Port25, and we're delighted to be here today to talk about our experiences and perspectives of adopting these new protocols and standards.

A quick background, Port25 is, as many people probably recognize the TCP Port, Port25 but maybe not the company, and our background and what we're best known for is a product by the name of Power MTA. We are an email infrastructure company so commercial MTA provider, and really our focus has been the community of legitimate senders and providing a solution that meets the specific needs around CRM, email marketing and customer communications.

So some of our customers include some of the leading email service providers. About 20 percent of the Email Service Provider Coalition are customers of ours, along with many of the large consumer brands such as Bank of America and Travelocity and Mary Kay Cosmetics and others.

In addition to serving that market, we also have another version of our product which can be deployed as an embedded component, for example, in an email security solution as an alternative to an open source component as well, and really what we see as our opportunity and mission is the adoption of the email practices that we're discussing here, and certainly authentication is the first one.

But really it's the beginning of a whole road map of new paradigms and certainly a great opportunity for email, but also a changing of the email infrastructure. This isn't going to be a point in time issue. This is really the beginning of an overall evolution.

So the perspective we want to speak to is certainly where we've been focusing, again enabling legitimate senders, and what are we hearing and seeing in the market from those players? And I think the reality is it's kind of a mixed message.

On one hand you have a lot of questions out there, and certainly in the noise of what's occurred over the last 12 months, there's been some confusion, and a lot of the folks we talk to are confused. The very good news is that they are still moving forward and certainly that speaks to the fact that senders are really incented to adopt these technologies.

Anything that a legitimate sender can do to help separate the wheat from the chaff they're going to want to do, and certainly in the noise of the market, what has bubbled up and what we were hearing that people are moving forward with is SPF, Sender ID and DomainKeys, and my little figure there is running.

Certainly everyone is not running at the same speed of course. We certainly find the email service providers actually are doing a great job, which again is probably not a big surprise. I spoke to Trevor Hughes in the hall, Chairman of the ESPC today, and he said as far as he's aware, every email service provider has published SPF records, at least SPF version 1, and that's a real credit to the group there and the focus that that community has.

Certainly since they're in the business of delivering email, it behooves them to move quickly on these things. Certainly large enterprises don't have

the same luxury. While they are trying to move forward quickly, what we find is as with any big corporate IT issue, a DNS change for example can take 30 to 60 days so your ability to move quickly and respond to issues certainly is going to be inhibited if that's the environment that you're working in.

When we think about the challenges ahead, if that's what's happening today in the market, the challenges ahead, the big risk is not that we can't figure out any point technology. It's really that there are so many new things that are being ejected that the complexity gets overwhelming, and I think that's as a community something that we need to be mindful of as we think about the battling standards, to make sure we're not expecting too much as far as adoption.

And so it is the various standards and the various versions that they're going to undergo and have undergone and there's all the different elements that have to be coordinated to make those standards work, and then there's a whole life cycle associated with managing those things.

So at times we get focused on the algorithm or the specifics of the technology. If we step back like any IT element that's dropped into an enterprise, it's really managing over time which is where the real cost

is.

And so when we think about helping centers deal with adopting these tools, while there certainly is I think a valid perspective that the IP schemes are rather straightforward in terms of their requiring fundamentally no DNS change, there's a whole life cycle of those managing those that is a little more complex, so in September we rolled out our first version of these products, and it supported SPF, Sender ID and DomainKeys, and what the tool set did was certainly as an MTA, we were very focused on the execute functions or the ability to stamp messages, for example, with DomainKeys.

But we also were really mindful of what it takes to configure a policy that is consistent with what your organization wants to do, and we're mindfully of providing a set of tools so that folks can send test messages, for example, to validate that, Oh, yes, I am complying with the standard, and then there's an ongoing opportunity to monitor that in fact your DNS is still linked up with your infrastructure.

We have ability to, for example, put the server in a mode so that it would not allow forged emails or emails that don't comply with the various standards to be sent through the server, and those are examples of how we're going to we hope help people manage the

complexity associated with adopting these standards.

So certainly one of the big focuses of this panel is testing, and as we've thought about the testing, certainly it begins with the functional test at a product level, and make sure that we're conforming with the specifications and the white box and black box test that you would expect, and then we go from there to the operational testing which addresses issues like performance and so forth.

I think the good news is that a lot of our bench marketing data, particularly as it relates to DomainKeys and the crypto approaches, corroborates with what we've seen Sendmail, the data that they published and also ColdSpark, you mentioned particularly with small keys, that the CPU utilization is not a huge problem.

One of the things we have seen, however, is as the key sizes get bigger, as you would expect, then the CPU problem can very well become a real bottleneck, and if you would go from a key size, let's say five twelve bytes up to ten, twenty-four, now you're talking about maybe a 20 percent hit on CPU going from a 80 to a 90 percent hit on CPU, and the resulting impact of throughput with the larger keys is in fact very significant.

So I think as we continue to test and evolve

these, I think we have to be mindful of the exact parameters we're using in the test. I know Sendmail testing has been great out their as a benchmark based on 384 bit key, which is actually below what the current spec calls for as a five-twelve bit key, and we don't think that will be material, but we think it's a scenario where we're going to continue to test and evaluate and hopefully collaborate with some of our peers here.

So we feel like we've made some good progress in terms of what we can do within the company. Where we feel like there is plenty of work to do is figure out how to make sure that implementations are in fact intra operable with other implementations, and I think that applies at a functional level as well as at a performance level.

And when we kind of have all those boxes checked off is really when we're going to feel very confident as it relates to consumer readiness.

So finally I think we just wanted to quickly close with being a bit I guess prescriptive about what we see some of the opportunities are as a community coming out of the this event and so forth, and I think as again we talk to customers, the issue of communication and having some clear message about where

we're going and deployment time frames, I think whatever we can do to help make that clear is really going to help.

I think certainly inside the industry we're still in the process of building the ship, and to the degree we can make the communications clear to everyone, I think that would certainly help.

In addition, as we move forward with the various proposals, to make sure that to the degree possible we can have standard test beds that folks can rely on. I know right now there is probably, I don't know 10 or 15 different testing servers that people have set up that can receive different messages and tell you whether or not it's going to work.

We have one, and there's plenty of others out there. It would be really confidence building if we had maybe one that we all had a great deal of confidence in where we could send a series of these and other to battery of test data to help.

I think one of the other key points is that I think we need to be mindful that as much as this is a technology challenge, it's perhaps more so a marketing challenge, and in that we're asking an entire world really to adopt the new technology and what is necessary to effectively accomplish that in a marketing level is

really significant, and making it clear why we have multiple standards and how they work together.

I think that is one example of -- maybe folks in this room understand that but in the broad market I don't think that's very well understood. Taking those messages and making them clear I think is important.

Since as everyone has said the authentication is only a building block, we're asking folks to do a lot of work, but the reality is the standard problem is not going to be solved overnight and making sure that the expectation is aligned there I think is key.

While I think the grass roots bottoms up efforts are really important, perhaps as we get further along, having some type of campaign, maybe analogous to a Got Milk campaign for email authentication, maybe we should think about these things as well to help really hit the broad market with what the effort is and what the benefit is and a concept that we have batted around is perhaps it's installed under ID required, and if you go through a certain certification process, you can put on your web site that, yes, you have been certified in the various flavors of email authentication, and that becomes part of how you talk about yourself as well.

Thank you.

(Applause.)

MS. COLEMAN: That was really great. You know, so far I've heard a lot of conflicting information. I've been taking notes as you all are as well. So far we've had one panelist tell us, "It's time to deploy." We had another panelist say he won't even go there with respect to where we are in terms of implementing these, so we're from one extreme to the other.

So let's hear from more of our remaining panelists. Maybe we can reach some consensus about this by the end. Let's see.

Now we have Barry Leiba.

MR. LEIBA: Hi. I wasn't going to go through this item but Sana said we had to entertain you, so I'll start by entertaining you with a little fact that will probably surprise some of you, and some of you have been around long enough to know it.

I'll go back to one of David Fowlers's charts where he had this sleeping arrow that started on the left of the screen and moved to the right of the screen and had sort of different stages in email along there and what we used it for.

The question that I ask people is: Where on that time line, if you put a time line on that arrow, where would spam have started? Where was the first documentation publicly written about the spam problem?

And my answer to that, if you think about the screen, raise your hand, somewhere around there.

The late John Postel wrote an RFC for the IETF in 1975 about the spam problem. That's almost 30 years ago, so when people say we're not going to solve the spam problem overnight, well, yeah, because we've been working on it for 30 years.

So maybe that entertained you. Maybe it didn't. I'm afraid my accent isn't as funny as David's so I can't do that.

MR. FOWLER: Ouch.

MR. LEIBA: So anyway, I'll start with goals that we have for the various sender verification things, and this conference is called Email Authentication, but I've switched terms, and I'm calling it sender verification, and maybe I should change that even a little and call it sender validation, because what I think we're really trying to do here is not to actually do hard authentication like we would when you log on, but to a great extent what we're really trying to do is determine with a reasonable degree of certainty where the message came from or at least that it came from where it said it came from.

In this case a hundred percent authentication isn't necessary. We're trying to attack the problem

reduce the problem, and on all my slides, you're going to see reduce, improve, those sorts of words. We're not claiming that we can solve the problem. Only that we can make it better.

So we're going to increase the efficacy of other mechanisms that we have. We have whitelists and blacklists now which I'll call good and bad sender lists on my charts, and having a better idea of where the message came from makes those more effective. For legal efforts it helps to track down people if we have a better idea of where it did or didn't get from.

For challenge response systems, we're challenging mailing lists and robots, now challenged responses have become joe-jobs now, just like bounces, because we're challenging the wrong entity. This will help that. Phishing obviously we're trying to attack, and we've said a lot about bad bounces, joe-jobs.

I've showed this chart a lot. To the left we have the legal action that we can take against spam. On the right we have this hierarchy of technical mechanisms so we have challenge response systems. We have identification of where the mail came from, payments, whitelists, blacklists, content analysis.

We also have got the personal preferences here, and I'll go back to the previous speaker and agree that

it's very important actually I think it was the second Karl that said that personal preferences were an important piece of this, every user is going to have a different view how they want their spam treated.

Within IBM we've had some groups who insisted that just for the internal, the mail that we get internally, we've had some groups that insist they may not get any pornographic spam and they don't care what the false positive rate is that it takes to achieve that.

We've had other groups particularly marketing groups who need to get mail from customers who said we can't have false positives and that if that means a little porn gets through, we'll delete it, and you can't make both of those happen at the same time without having some sort of personal and organizational preferences involved there.

Identity is what we're talking about here, making some sort of identification of where it came from, and I've just got the little -- my animation here that shows what it enables. It enables all these other things and probably also has something to do with content analysis.

This is similar to a chart we've seen earlier today, so I won't go into it a lot. I'm comparing IP

address based mechanisms with signature based mechanisms, and let me quickly look over it and see if there's something that hasn't already been said.

Basically the different points of the transmission where it works, whether the message being modified along the way affects it, how well it can deal with forwarding. The layering is interesting. The IP address mechanism, this IP address is authorized or isn't with signatures we could, if we set it up that way, have multiple layers of signatures on the message and validate several pieces along the way.

Simplicity of implementation, DNS, okay. The one, the signature, can use public key infrastructure, we've punted on that as I had a discussion back here with the people from NIST about how we've not been able to solve public key infrastructure, but if we ever do, we have that there.

I'll skip the rest of this and go to limitations. With any of these, we have to be very careful about what we say we're going to validate, and we're only going to validate what we say we are. This is not a -- this has been said. It's not something that -- I'm sorry, I lost my train of thought.

We have several different mechanisms, several different fields that say where the message came from,

and we have to be very careful about what we say we validate compared to what we actually are validating.

In many cases we've seen people who said the spammers are signing up for SPF, are publishing SPF records. The spammers and phishers simply admit who they are to the infrastructure, but what does the user see, and the user still sees the spam or still sees the phishing attempt.

If the spamming domain doesn't participate, we can only say that that means we put it through some more filters, some more careful scrutiny. AOL has said that they're not willing to delete mail based on the lack of these, so it's important for the legitimate domains to participate so we can whitelist them or treat them with less suspicion. It's not sufficient though.

It's still possible to control the end users, and I agree with what Dave Kaefer said earlier today about in principle, we can't require changes to the user interface to enable all of this, but in practice, looking at what the ISPs are saying about not being willing to trust just what happens here, we've got to have changes to the user interface to show the user what is and isn't to be trusted, that's especially true with phishing.

So to the purpose of this, testing. We're

focusing on what we need to test, and I thought it was very cool that the first one we had showed some numbers. Now, I'm not going to show you any numbers. What I'm going to talk about is some things that we have to be careful that we do test as we go through this.

We have to test how these systems work with legitimate senders that don't participate in the system we're doing. That's sort of obvious. The other side is we have to test with how we deal with spammers who do participate and phishers who do participate. Can these systems still be effective against those people?

We have to test it with transient failures, what appears if a DNS lookup fails temporarily, and we have to test against non transient failures, what happens when we go through a forwarder or a list server that modifies the header, modifies the body.

We have to test with anonymous mail, and we have to make sure that whatever do allows anonymous mail. I'll go back to the first thing this morning where we had quite a discussion about that. IBM strongly believes we need to make sure that whatever we do still allows anonymous mail and free speech.

Finally, can this be used as evidence in court, an issue that I can't answer but something that the lawyers have to consider as we go through these

proposals and we go through testing them.

The final thing I'll say is that there's no answer to the spam problem. We aren't going to solve the spam problem. I believe we're not going to solve the spam problem ever, that there will always be spam. We want to keep it under control so that email is still usable and people can still trust what happens with it.

So what we need to do is have as many approaches to it as possible. SPF is there. SPF/Sender ID is there. CSV is there. DomainKeys and IIM are there, and there are many other mechanisms that we're all talking about, and we're all using and we have to use them together. Most of them help. Most of them also cause additional problems. Perhaps several of them together can mitigate each other's problems and give us a better answer.

The other thing is that everybody -- each of the mechanisms has its fanatical supporters, and I don't mean that in a negative way, but what we do have to be careful is that that doesn't lead us down a path that we favor one and forgo all of the others. We have to merge them. We have to use them together.

Finally open standards, open standards, open standards. That's what we're all here -- that's what a lot of us from the IETF have to work with. Yeah, I got

some laughs over here. Okay. Anyway that's the end for me.

(Applause.)

MS. COLEMAN: Thanks a lot, Barry. We appreciate that. I think that you've raised some good questions there about kind of standardizing in a sense what we're testing for, and one of our earlier panelists, I think it was Bill, said there is no uniform testing methodology, so these are all things we can think about.

We're saying we're doing testing, but does it really mean anything if we're all doing our own thing coming up with different results? So with that in mind we'll give the floor to Dan Nadir.

MR. NADIR: Thank you. I just want to echo probably most of what Barry just said. He said a lot of it more eloquently than I probably will. FrontBridge is a managed service provider for anti-spam, anti-virus, stuff like that, so people change their MX records. Mail flows through us and we deliver it, so really we're consumers of all of this technology.

We don't really care. If it works, if it adds good value, and if it doesn't break anything, then we're inclined to want to do it. Early on we were looking at SPF and I'll say /Sender ID now. For us it was all

about ease of use. It was easy to do, and we predict that people will be more likely to do it because it's easy to do or it's relatively simple.

And we don't have sort of -- we have low expectations, let's put it that way, right? We're not looking for something that's going to fix everything right away. We hear a lot of arguments and someone will say, "Oh, I have this great technology" and someone else will say, "Well, that will never work because there's one case out of a million where someone could do this," and then you're totally screwed. So we'll sort of accept that, but if it adds value and it doesn't break anything, we're likely to do it.

For us the interest was really and is really in phishing scams as much as it was for spams. So we have a spam filter. It works decently. We're not actually convinced it's going to do a great job in helping us prevent a lot of spam, but it does seem pretty clear that you can do better authentication. You're going to do better job of blocking some phishing scams.

We have relatively small samples so my data is not great, but we're finding that there's a lot of legitimate domains that are doing SPF. There are a lot of spammer domains doing SPF. It isn't clear that that's going to help us very much at all.

We were surprised that none of the big phishing targets are doing SPF or Sender ID. Again I recognize that it isn't perfect, and there are probably lots of different reasons, and people will say, well, SPF doesn't really help, Sender ID might help a little bit better but someone could do EBay-Billing, and they're going to get around it anyway.

We're willing to take that if we can make sure that something isn't coming from EBay.com, that it's coming from EBay-billing.com, that's okay with us. We're still better off than we were before.

Like everybody else, we're not actually blocking email. We're kind of experimenting with our rules and how we tweak what we see based on the SPF record, so we can tune our score one way or another. About 4 percent, maybe it's 5 percent of email today actually has an SPF record, and I have to admit I'm not really sure if that's the number of domains we see or the volume of email.

In general it's fairly low, and we do know for sure that we're not going to, in my lifetime probably or at least in my short career lifetime here, absolutely block or absolutely allow either way just based on the SPF record. We may allow based on some other things that we know about in an organization, but we don't

believe SPF is going to be the thing that we use for lots of authentication in general.

We do believe that over time it's going to help with fighting spam, but again just like everybody else, it's all about it's about reputation, it's about accreditation, so it's about knowing much more about an IP or a domain than just whether it passed an authentication check.

We think in the short term whitelisting is going to be a good idea and you have to just do it. There are probably going to be organizations that aren't doing the right thing with records, but we're still going to want to let their mail through.

Our customers are all business customers. We don't have any consumer users so it's important that they get their mail, and anything that would block any legitimate mail for us is really, really bad, so we don't want to do that, and if we can sort of work manually in sort of Jerry-rigging a system to allow some mail through, that's what we're likely to do.

Scaleability is key. We don't really have a good sense yet for how this is going to work when lots of people are doing this, and Barry mentioned things like DNS time-outs and what's going to happen when we're processing 150 million messages a day, and there start

to be errors or people are not configuring things, so it just isn't clear to us that it isn't going to scale, but we hope so.

There are a lot of edge cases, and we don't know what we don't know, and it's kind of scary. That's why I think, we're as AOL is doing, sort of taking very careful steps. We want to balance the really, really edge cases that might break again where it's affecting only a couple of people versus sort of these weird edge cases like mobile phone, email, where we just can't block or we can't make decisions based on some kind of oddity.

We're also seeing that there's variances in configuration. Like someone was telling me that our customers are getting confused about, do they do a redirect?, do they do an include? It's not clear. They're confused so we have to help them. The nice thing about it is for our customers, it's a one line entry. We don't have to really do much. We can do that for them and everything will pretty much work.

We still don't know what to tell them about the future of Sender ID and what's been happening or what they should do, but we're monitoring it really closely, and we do think that there's a lot of I'll call it pseudo good email that people are considering sending.

Every time I get something that says it's from a friend of mine, I open it up, and it says, "Bob thought you might like this newsletter or something," and I go, "Okay, that's great, I'm not going to get that." There's a lot of email that's getting forwarded around. That stuff we think isn't going to work, and people are going to have to either change the way they do it or people like us are going to have to make some decisions about how we treat that kind of email.

Again we're all about being pragmatic. If it helps us, and it is overall going to be better than what we have today because most of this stuff is better than what we have today which is like nothing, so if we can do something and it helps us, we're in favor of it, so that's what we would like to do.

So I just pretty much said this, right? Are we still excited? Absolutely. We don't think it solves the problem. We don't think it's going to solve the problem. That's not what we're after. We're after data. It's just a better data point for us. If we can get to the point where we have sort of the high road and the low road, the high road we don't really apply a whole lot of additional checks to, and it's much more likely the email is going to get through, and we've got the low road where we apply a lot of aggressive checks,

and it's much less likely that email is going to get through, then we believe we will have succeeded and again we'll be better off than we are today.

That's it.

(Applause.)

MS. COLEMAN: That's great. Thanks, Dan. We appreciate that.

Now we're going to hear from Robert Sanders. You can feel free to come up and provide some remarks, no visuals required remarks.

MR. SCHNELL: I did not come bearing slides.

MS. COLEMAN: We won't hold it against you.

MR. SANDERS: Can everyone here me okay? Great. So there's been a lot of cautious optimism about authentication of emails so far, and I came prepared to echo the same, but I think we need some balance, so I'm going to switch it around a little bit and provide some perspective from a consumer ISP that also actually does a fair amount of business service and has a slightly different take on things.

So EarthLink has about 300,000 domains we manage for businesses, about 140 consumer domains, so we have a somewhat different perspective from say AOL who has, as Carl said, a very small number. We have a user base that is very heterogenous. They are not web based all

together. Many are. They are not using a single email client. They are all using various POP 3 and IMAP clients and SMTP clients to send mail through us. These clients have been configured in various different ways. Some of them don't provide us with any authentication at all and we use just the IP address that they're sending from to allow service.

In some cases they do provide authentication, and we are certainly working to get more of that, but the basic point is we don't know enough about our customers and who is sending mail to be able to really provide much authentication information to the recipients of the email, and it is an expensive process for us to get to that point, and so the question is obviously: Is it worth it? I think today without reputation, from what we've seen so far it probably is not.

So we have some interesting numbers that have just come out of a study that we've done over a fairly small corpus, about a hundred million messages I would say, and I should say that these measurements, as has been mentioned, are not easy to interpret from one period to the next because we all have very different email systems, very different customers and very different spam filtering systems.

But from the mail that we do see, from the domains that have SPF records published, about 90 percent of the mail that passes SPF is spam. 90 percent of the mail that fails SPF verification is spam, and so forth, down through all the various SPF result codes. You can interpret that various different ways.

What's interesting is for domains not publishing SPF, only 40 percent of the mail we received is spam, so for us the primary purpose of SPF records is a great spam sign. You can also say that argues for the efficacy of our other spam filters, and I will certainly take this, but it is interesting.

Why do this at all, and I think with reputation, we can do a lot of things with this, but the idea that we'll get something out of it for a little while until the reputation comes along, I think that's already been disproven for a lot of us. So maybe it's not going to stop spam.

What is it going to do about phishing? Well, as has been pointed out many times, you may not be able to send a Citibank.com, but you'll send a Citibank-Accounts.com if you're a smart phisher and users will not be able to tell the difference, and there's no way we'll know the difference as someone who is receiving that email.

They are who they claim to be. We don't know if they are who they appear to be, and that's why I would echo what Barry and others have said. There has to be some consideration of not just how to feed this data into filtering algorithms, but how to present it to the user and let him make an informed choice about it.

We actually have a tool called Spam Blocker which we have deployed to anyone who wants to download it, and its purpose is to say well, we don't control all the email they get. In fact many of the users are not our customers, though we can control the web sites they go to, and so we basically have an ad hoc reputation system using URLs fed to us from Brightmail and EBay and various other partners.

That has actually been very successful in preventing phishing success with our customers. Some of the numbers I have here I find kind of interesting. As of last year, a phishing attack on our customer base cost us around $100,000 just in terms of call center impact, and that was around 20,000 calls per incident. We are down to about 300 calls about $4,700 per incident, and that's without really I think changing much of how we filtered the mail.

So again if authentication is not going to stop spam and it's not going to stop phishing or we have

other tools to do so, is it worth the investment? And I'll tell you why it's an investment issue for us and also why I'm a little bit afraid of what both authentication and in fact certain kinds of reputation might due to affect an ISP like us.

So reputation hasn't really, really been well defined, and that's on purpose. It's out of scope of many of the things we've done. Think of reputation as a function over something mapping to something, in this case generally it's assumed over a domain or a sending host and returning some value which generally also hasn't been defined, but let's call it probability that a message from that domain is spam, which is a useful thing to have.

I don't know whether that's the only useful reputation function, and I think it's more useful to some domains than others or more tolerable. From an ECommerce site, which is a very heterogenous type system, Amazon, for example, the reputation function is generally going to be a very useful thing, because generally if the mail is actually from Amazon and SPF or DomainKeys or whatever will give you that, then generally the mail will more or less be sent legitimately from a small controlled set of people.

However, reputation function applied to a domain

like Earthlink which has tens of millions of mail boxes doesn't give you anything interesting. A lot of users are good. Some of the users are bad. Everyone gets the average, and so that's troubling.

It would sort of -- contrary to what many have believed, which is that these systems will shut out the small business or the small ISP, it actually makes it very hard for a large ISP to maintain a useful reputation. That to me says we need to think about reputation down to the user level or accreditation from the ISP to perhaps even a current message basis to say, I have a high confidence that this message is spam or is not spam hopefully as opposed to just EarthLink.Net has a pretty good track record of blocking spam. That's a very difficult thing for us to really motivate around.

So what are we actually doing though? Like I said, I came prepared to be optimistic, and I really am deep down. We are publishing SPF V 1 records and will continue to do so and then upgrades. We are not blocking mail based on SPF failure. In fact, we're not even really verifying it in real time. We're just post processing. We are planning to implement DomainKeys.

That's actually in process right now using the generously provided open source library. We are very impressed with some parts of IIM, I guess the parts --

the additional parts and would love to see those two merge, and certainly would prefer to have only one signing scheme to test.

It's not likely that we're going to sign a message twice. We may publish two different kinds of IP records, but we're not going to double sign.

We certainly have seen that our practices, like Port25 blocking, actually make some of these systems more difficult to support. If the user cannot connect back home to his authorized mail server, then he can't really benefit from these authentication schemes, not the IP addressed based ones certainly and not the cryptographic ones without user keys, so we have -- although we do Port25 blocking, we have deployed Port 587 as a submission Port so that our traveling users can get back to us, and we highly encourage others to do the same.

Port25 blocking, although it does make authentication more difficult to deploy, from our point of view is a responsible thing for an ISP to do, and we think it has actually stopped a lot of spam.

We are, as I said, converting our user base to a more strongly authenticated configuration where we can, although with zombies and Trojans I'm not sure how much that's worth. Once we assign more value to the user

credentials, they will get stolen more often, and I think that maybe suggests that we should look at other ways of controlling access to the system.

People have even suggested two factor authentication. In fact I think AOL is currently selling that and congratulations, Carl, very prescient move.

That's not the only way. I mean, certainly you can limit the value of the credentials by rate limiting as we are doing and others do as well, but certainly I think that the zombie problem has tossed a lot of this on its side, and we're going to be doing outbound signing where we can.

We are in a sense doing SPF where we can, but we are doing it in a way that many domains are doing it, which is to says these are our mail servers but you can get email really from anywhere else, and it's still valid.

I think it's very difficult for an ISP to take that last caveat away, an ISP of our sort, but we would love to get there and certainly will as soon as we can.

Most importantly I think we are going to be sharing this test data and have already started to do so within what's called MAAWG, the Messaging Anti-Abuse Working Group. I would encourage everyone that has this

sort of data to get involved there. I think it's going to be difficult to share certain kinds of data, in particular things like per message failure or success, for some of the cryptographic schemes to see are they or are they really not working end to end, but general statistical data I think we could collect there.

And we'll be updating our systems including user interfaces for users, including clients and so forth to support and display, to present to the user authentication information and hopefully reputation as soon as it is available.

I believe that's all.

(Applause.)

MS. COLEMAN: That was great. Thanks, Robert. I think you touched on a lot of key points there, particularly your last point about sharing information in the MAAWG forum perhaps and in other locations where we can get a sense of what we're all coming up with, compare how we came up with it and move forward from there, so we appreciate that. What you lacked in visuals, you certainly made up for, and we appreciate that.

Now we have Ron Schnell from Equifax.

MR. SCHNELL: Thank you. Equifax, founded in the 1800s as a company that gathered and published

information about the paying habits of retail store customers. Today, we're the leading provider of data information for consumer initiated transactions.

We host the largest and most comprehensive network of automated consumer credit information in the U.S. and Canada, and we have over 300,000 customers that use us to evaluate risk, protect against identity fraud and market products and services.

So why is Equifax interested in email Authentication? Number one, we're concerned about the future of email, as its usefulness may be declining due to spam. We have a great interest in the financial sector, of course, and we feel that phishing is a real concern for us and our largest customers, and we're a technology company with strong expertise in identity protection and verification. After all, we're one of the earliest reputation services. We've been doing it for 105 years, and delivery of email to our consumers is of vital importance to our business.

So our thought process in trying to implement and test these methods, phishing came first, and we started to think, Is this going to help the phishing problem. Phishing of course is easy because email was designed with no authentication in mind.

When I started on the ARPANET, the early days of

SMTP, everyone who was on the ARPANET knew each other by name, sometimes by face so it was never even a thought that we needed to authenticate email.

Although the era of trust on the Internet was gone a long time ago, it's still the case that the mind's first instinct is to trust what the email client is telling you, so even me, I've been in this industry for awhile, when I see emails from financial institution customer service, my very first instinct is, "Oh, this email is from this bank," and then half a millisecond later I say, "Wait, that's probably not true," but still the first instinct is always when you're looking at it that it must be true.

So one thing, I've been talking to people even inside my company, and there's sort of this misconception that with email authentication, phishing is going to be helped in some way. I hear people say, well, all that you need is you need to have the banks and the people that are targets of phishing publish sender ID records with PRA, and then if it's not the correct -- if it's not who it purports to be, then it won't get in the inbox.

But in reality people who are duped by phishing scams -- we know they're not duped from the Mail From. They're also not duped by the purported sender because

most email clients don't show even what's in the RFC 2822 from address. They're just showing the pretty name, and the pretty name is what's in quotes after the email address, so Ron Schnell or Citibank customer service in quotes.

It's not going to show even account service at Citibank.com or some people have suggested, they'll just get around it by adding a dash or something else. It doesn't matter because it's really just the stuff in quotes that most email clients show and some web based email shows, so really it's not that simple. It's more than just the targets of phishing that would need to subscribe to one of these methods to help the phishing problem.

So if you help the spam problem, however, you will help the phishing problem. Spam is also driven by the ability to send email without authentication, and widespread adoption by email providers and sending domains would be required to have a chance at a measurable effect on spam. So that's something that's going to take awhile.

Now, once we have widespread adoption, it becomes a useful enforcement tool for agencies fighting spam like the FTC so that's probably one of the things that's driving this Summit. Until widespread

implementation by email providers, unless unauthenticated email is rejected out of hand, authentication is not enough to help spam. We've heard that a number times today so I won't dwell on it.

But if only authenticated email is allowed in the inbox, useful decisions about email can be put in the hands of the end user, and a few people on this panel have talked about that. I think it's a great idea. The only way you could really do it though is if you were to throw out all the email that didn't subscribe to a method of authentication, and no one here has really suggested that, but it's interesting to look at what you could do if that happened.

You could have meaningful user maintained whitelists, meaningful user maintained blacklists, and something that power email users have been hoping for for a long time, automatic folder management, so that if you know an email is from a certain person, you can put it in a certain folder. Right now you can't trust that.

Now, also if you did that, the privacy concerns that were brought up early this morning would come into play, and I don't think those should be minimized. We should certainly keep that in mind when you're setting your user policy based on what to do, based on whether the email is authenticated or who the person is, but I

think that should also be put in the hands of the individual user.

To address Paula's political free speech concern from this morning, perhaps government entities shouldn't be allowed to just throw out unauthenticated email. That's one way to get around that.

Talking about user maintained whitelists because it's sort of a favorite topic of mine, if users only allow email from senders from whom they expect to receive communications, this would greatly reduce the spam problem, but of course what that does is it changes the way people use email. Everybody's been used to email being open for the last 25, 30 years, and our society is not ready to address a drastic change like that to email or so it seems. This is more similar to the way people use Instant Messenger which has grown at an incredible pace.

So you can set up your Instant Messenger so that you'll only receive messages from people from whom you're expecting to receive them, so it's interesting that people will accept that from Instant Messenger but not from email, so it's probably just a matter of history and the way people are trained.

So one thing I think we could do, if we wanted to make a more restrictive email, is just describe it as

we're actually enhancing Instant Messenger and we're adding email features to Instant Messenger and then you'll end up with email that has that authentication just like Instant Messenger already has, and maybe people would be willing to accept it.

What people seem to be afraid of here is email is going to go down the tubes and it's not going to be useful anymore, and I argue it's barely useful now, but what's the alternative? The alternative may be to enhance Instant Messenger, make that the business email, add storing power and make it store messages and use that for your first class email and leave the old email for a third class email. That's just a suggestion I like to get out.

So I'll add again, like everyone else, that reputation services are an important adjunct to sender authentication. Users will need help in deciding from whom they want to receive commercial email, and reputation services are probably the best tool.

Some users will still rely on their email provider to make the decision for them. Maybe they don't want to. Maybe they don't understand it well enough, or maybe because authentication isn't widely implemented enough, and email providers' use of reputation services can really help with that.

So what happened when Equifax decided to try to implement some form of authentication can be described pretty easily. We began following Caller ID, and George Webb at Microsoft was kind enough to ask for our opinion on that, and we gave him some notes.

We started looking at DomainKeys, and then all of a sudden out of nowhere SPF immediately became the front runner for us for three reasons: Easy implementation, seemed to be having wide Internet community acceptance, but then most importantly, AOL made a statement, "If you're not using SPF, you're not on the whitelist anymore."

So although SPF is not necessarily a solution to spam or phishing on its own, for us implementation became necessary to ensure delivery of our transactional and marketing messages, which goes right to our bottom line.

So we found that mass confusion surrounding the various proposals existed. Issues including intellectual property, privacy, obstinateness, which may be a strong term, but I'm not talking about today. I'm talking about a long time ago, like a week and a half or so.

Once we got past the problem of which methods to test, numerous implementation issues arose. Because

Equifax acts as a transactional mailer, a marketing mailer and in some cases an email service provider. Which SPF records to publish is not straightforward, especially with PRA requirements looming.

For email service providers, it is particularly confusing, who is the responsible address and who should be on the envelope? I subscribe to the SRS discussion. There's a great article by John Glube, who talks about the perspective of an email service provider, and there are about eight different possibilities that you should put for each of these, and no one really knew the right answers. There were some suggestions, maybe you should do this or maybe you should do that but there was never really a consensus.

As it is right now, SPF 1 technical implementation is quite easy, and it went quite smoothly for us. All our transactional marketing domains now have SPF 1 records published. Pretty much the only test result we have to give you is that Gmail successfully recognizes our SPF records and adds little tags so we're happy about that, but there's no recognizable improvement in our deliverability or ISP relationships that can be attributed to our publishing SPF records.

We did subscribe to a reputation accreditation service for our outbound mail. We had mixed results so

we're not subscribing to that anymore, and we could not find an SPF plug into Lotus Domino for our corporate email, so I have no testing results to give you for how it affects spam coming inbound, but from what I hear it's a pretty low percentage anyway.

So in summary, implementation of our chosen email authentication method was easy to perform on the sending side but no benefits can be appreciated until wide scale adoption takes place. Our selection of the chosen method was not based upon scientific merit but had to be based upon our business critical needs, which was based upon the opinion of the largest email providers.

The current state of flux and confusion surrounding the major proposals are such that it would not be prudent to spend a lot of money to implement right now. It seems to be changing. I think this Summit is probably going to be helpful with that, and we're certainly going to keep an eye on it, so I look forward to your questions.

Thank you.

(Applause.)

MS. COLEMAN: All right. Rand Wacker, come on down, our final panelist, and following your presentation we'll take questions from you all.

MR. WACKER: Thank you very much. My name is Rand Wacker, and I work for Sendmail, which is a hybrid open source and commercial company providing email solutions to Global 2000 enterprises, ISPs and also a wide array of small senders who are using the free version of the MTA that's been available for more than 20 years.

We have been working with a number of authentication proposals for the past 12 to 18 months and we've implemented and released it for testing open source versions of DomainKeys, SPF and Sender ID.

Now, having been on the World Cup tour with many of these folks for this past year, I have to say I agree with most everything they've said, and we've had similar results to what they've gone over, so instead of kind of rehashing some of the similar numbers, I wanted to talk about some of our testing results from an implementation standpoint of our customers and what our recommendations are for people right now moving forward.

So some of the things that are interesting about these proposals are not necessarily the technical aspects of the specifications themselves, but the changes to the business processes and the changes to the network architectures that people are going to have to do in order to enable authentication.

EarthLink has talked about some of the issues they're having, authenticating their end users before they relay mail through, some of the issues about Port25 blocking and enabling the submission port and whatnot, so it's important to know that roll out is not just a matter of putting some records in, and it's not just a matter of putting some software in.

A lot of effort is going to have to go into auditing your network and determining kind of what your business practices are for outbound email, be it from your corporate servers or remote users or third-party mailers who are currently sending mail on your behalf and who you want to authorize as well.

So we're recommending that people go through these processes because that kind of work is going to be the same amount of work you're going to have to do for all of the different proposals you roll out, and we recommend people roll out both IP and crypto based solutions because as a sender you're not able to detect which path a message may take based on recipient action.

So as a sender, you send a message to an address, you don't know the recipient has an address to be forwarded by a mail list or a forwarding service, and so your best bet as a sender is to as give as much information of the recipient as you can, so that includes an IP and a

crypto based solution.

Performance. We're seeing the same numbers on performance as everyone else. The bottom line is we're not really concerned about some of the overhead there.

I think where some of the recommendations get most interesting are what the receiver actually does with this information. We are recommending that people check multiple authentication methods and receivers be aware that most of the time that a receiver fails authentication, assuming that the record published wasn't broken or if the signature was applied properly when it was sent out, most of the time, when a legitimate message fails authentication, it's because of an action the receiver requested, be it forwarding or be it some interesting path that the message went through.

So we're in a transitional state where we're looking at a time when receivers should be comparing the results of their authentication against the classical spam scanning they have now. By looking at a message that may have failed an authentication check but would have otherwise been considered to not be spam, then that's a good way to ferret out the broken forwarders and the paths that they're going to need to be able to fix in order to make this a true reliable authentication system in the future.

So what do you do with the authentication failure? You have to decide if you're going to reject something out of hand or possibly accept it as either unauthenticated or process it slightly harsher.

One of the things that we are recommending is that people do not necessarily discard email directly. We think that silent discards have made emails somewhat unreliable, and we want to see people actually rejecting the messages so there's a positive feedback to the sender. We need to get back to the point where if something goes wrong, you as a sender know something went wrong and you can fix it.

Finally, the question is what do you actually give to the end user? Some people have talked about the idea of the SSL lock or a gold star or a green light on the message coming in. Every different ISP, every different MUA is probably going to implement these in different ways. What we're recommending is people be gradual in rolling out these kinds of changes to the end users.

Maybe some of the things that they do first is that they strip off that pretty name that may not be able to authenticate or they only show it in the case of a known or trusted sender. What we want to be careful about is we don't want to start training or conditioning end

users to expect to see a green light or to accept broken authentication.

We want to see end users -- we want to see a lot of the work being done in the acceptance process on the server side and try to not leave the decisions up to the end users because it's confusing enough for all of us, and we don't necessarily want to push that confusion to the end users and just make the problem all that much worse.

So that's about all we have for now. Thank you very much for having us.

(Applause.)

MS. COLEMAN: Well, great. We've got folks out there with microphones. If you have questions, just put your hand up. There's one the gentleman in the white shirt.

MR. MESNIK: My name is Peter Mesnik, M-E-S-N-I-K. For those of you who have tested or have been testing the performance of the signed mail, what was the average size of the messages that you were using? What was the largest message size and did that have an effect on performance?

MS. COLEMAN: Okay. Scott?

MR. BROWN: I can talk to that first. Maybe not. So what we do is we do a distribution of message

size between 10k and 200k, weighted between 10 and 40k for the bulk of that mail to sort of simulate corporate mail with some spikes up.

It did have some impact. The bigger the message, the slower things are, the same for all things.

MS. COLEMAN: There's a follow-up question in the front here, if you could repeat that, sir.

MR. RITTER: My question was, was it different against the base line or was it proportional?

MR. BROWN: Yeah, it's different across the baseline across the board, so the bigger the message.

GEORGE RITTER: It doesn't matter?

MR. BROWN: It appears the majority of the work is in the SHA1 Hash.

MR. RITTER: Oh, George Ritter.

MS. COLEMAN: Oh, yes, let's have some more follow-up. Oh, was that Bill Karpovich?

MR. KARPOVICH: I was going to say our testing was similarly on an average message of 42k consistent with some of the tests that were published and was done as well, and clearly the size of the message does have an impact and as I mentioned, certainly also the size of the key that you use will have an impact on CPU utilization and throughput.

MS. COLEMAN: Great, great. Any other panelists who would like to respond? Okay. Let's take another question. This gentleman in the third row on the right.

MR. CHAFFEN: Steve Chaffen. I have a question. Only one of you I think really talked about zombies really, and I was told last week by somebody who works at HP in anti-spam that more than 50 percent of the spam comes from zombies.

Aren't you concerned about zombies suborning the reputation systems? I mean, if momandpop.com gets a good reputation, doesn't that make them a higher value target for someone to take over and then use their reputation or their credentials to send spam?

MS. COLEMAN: Who would like to respond?

MR. LEIBA: I have one thing to say about that. As my colleague from Earthlink said, they're blocking Port25 outbound, and that makes it -- that limits what the zombies can do. The zombies can't directly connect to outside SPF service.

MR. HUTZLER: Actually our experience, a lot of people talked about spammers registering domains and publishing SPF or Sender ID records for them. We've seen exactly the opposite with some of our fairly aggressive blocking or the zombies themselves. What

we've seen is the zombies are there, and the traffic is then routed through the SMTP servers at all the ISPs because most ISPs don't have authentication credentials required to send mail.

They trust the internal network so the architecture ends up being zombies as open relays. The main mail servers at the ISPs are forwarding servers for that traffic, and well over 80 percent of AOL spam today comes via that methodology. We see very little spam direct from zombies anymore. It gets routed.

MR. LEIBA: Doesn't rate limiting help there?

MR. HUTZLER: It should. Not many ISPs do that today. We're going to see that change happen hopefully over the next six months. But that is where we've seen the spammers go.

Perhaps if there is authentication on those servers, we'll see them go a different direction, but that's what we've seen.

MR. JACOB: I think to answer your question as well, it's a great example of why reputation systems are an important part of the defense against this kind of stuff because there are points that you can't trust the domain, and of course you've got to look at other of things and at that point just the content of that and the reputation of the content is important.

MS. COLEMAN: Great. We have one more. Let's start on this side. Let's see hands, please. Any questions on this side? There's a gentleman here, second row from the front.

MR. GILLUM: Hi, Elliot Gillum. Since we have this wonderful and diverse panel, we've talked about a number of times I think or we talked very close to it, a lot of different ways a lot of different times about spammers signing up for domain names, and nobody has really come out and said how much money the registrars are making off of all the domains names registered by the spammers.

I've heard rumors and rumblings about people upset about this, but do we have any concepts of what we might do to reign them back?

DR. BAKER: If I could, I would be glad to tell my shareholders that we are making a mentor out of this. The dirty little secret is a thing called a probationary period, and if you register a domain name and the registrar doesn't hand over the money instantly, if the credit card doesn't go through, they cannot pay for it. Most of those domain names that are used by the spammers are on stolen credit cards and cancel out very quickly.

So it's not really making anybody huge amounts

of money I don't believe. If it was the cost is coming out in other areas.

MS. COLEMAN: Any additional response from the panelists?

MR. CHADWICK: I think this is a key thing. The one thing we do is we focus very heavily on fraud protection, prevention, that kind of stuff because most people come in, spammers trying to buy domains are going to use a fraudulent credit card, and it's only going to be in the system for a couple hours before we catch it.

Not every registrar is as gung-ho as we are. We block orders, sometimes too many orders that creates problems to our customers, but there are so many registrars now, and there really are no real controls, that they can basically put their name up there, and they'll probably get it pretty quick and they can start sending email relatively quickly.

There is no 48 hour probationary period like that today. Basically once they buy the domain. They have the DNS entries, they can publish DNS right then and there depending on how DNS within a few hours depending on how DNS propagates their servers across the Internet, they can be sending spam.

I think there has to be better control at some point put into place during the purchasing process. The

transfer process, but that's going to take -- there are literally a ton of registrars now, and for one to do that kind of puts us outside the norm, and everyone must go through different registrars because it's easier to buy the name.

They're not worried about the fact that they're selling 5 percent of the names to spammers. They want to go where it's easy as possible and then get their domain in minutes and use it.

MR. HUTZLER: I would sort of add, I understand where you're coming from, and we've had this frustration at AOL for years. We used to block URLs by domain, still do, but a lot of them, and we would get frustrated seeing a spammer go through five, six, seven dollar domains at a thousand a clip, but I would sort of argue that it's a little indirect way to stop this.

You can even imagine. Gee, we'll have a blacklist and a white list for registrars, good ones and bad ones. We used to have the same problem with email service providers. They had clients that weren't the best clients in the world, and they had the same argument, rightfully so, that if they booted one of these huge clients off their network, who obviously was not sending legitimate mail, they would go to the next one, and we certainly saw that.

At the same time keep in mind that the spammers can easily high-jack a well-known ISP's domain and use that as their spam platform. We saw, not picking on Comcast, but in March they were in the news admitting they had a huge problem. I'm not going to name names, but there are plenty of people sitting in this room who we've had conversations with recently. Gmail is a great company. They're signing DomainKeys and we've received our first Gmail spam. Gmail is fantastic. They were on it instantly. They knew exactly what was going on.

But these technologies really are going to force the spammers to the ISPs. As the blacklisting companies, MAPS and so forth, hit really fast at blocking these zombies. There are only a few paths that are going to be left and/or if we implement these technologies, they're going to have to either buy domains, which is a very hard thing to do, especially if you have reputation, you also have to have or more than likely they're going to come sit on the ISP.

AOL has the same problem as every other ISP, and we have to combat it every day. In our case it's accounts that might get phished and then used as spam for a few hours, but we all have to really pay attention to this and look at our outbound problem as we go forward.

MS. COLEMAN: Great. Any more questions? Yes, you sir.

MR. HAMMER: Yes, Michael Hammer. Everybody's been talking about authentication schemes that are really, for the most part, domain name based. People like Dan Kaminski have shown that while interesting things you can do with DNS, are we just pushing the problem to a different area, that is, from one wide spread early protocol which has been resistant to change to another wide spread early developed protocol which may be resistant to changes of susceptible to subvergence?

MR. HUTZLER: I guess your question is sort of DNS's vulnerability and if we put a lot of stock in DNS, they'll compromise that?

MR. HAMMER: In other words, if DNS is susceptible, just how trustworthy are the authentication systems based on it?

MR. HUTZLER: Not that this explains it in a way, and I'm not an expert in DNS nor in ISP address and the ability to spoof a session, but those are two vulnerabilities you'll see named in I think almost every spec. Only as good as DNS is. If you can spoof your connecting IP address. We don't know how to attack that.

If you consider those two, those are doomsday scenarios, and if someone is able to spoof DNS on a large scale to enable them to spam, I think they could probably use it to route Amazon traffic to their own personal web page or Citibank traffic, let alone if they can spoof an IP they can be anybody on the Internet that they want.

If those core things get compromised on a large scale and all of a sudden became easy to do, like there was a scare about IP address sequencing that came about a few months ago, I think we're in a lot of trouble for the infrastructure of the Internet more than just email, so it probably is something you can count on.

MS. COLEMAN: Okay. Anybody else?

MR. SANDERS: I would like to comment on that. I will say that you're right, it probably is a doomsday scenario in regards to spam, but it takes just a few phisher messages to be successful to make it worthwhile, and that's why I think when we talk about fighting phishing with the systems, we should keep that in mind. The acceptable failure rate is much lower for a phishing solution.

MS. COLEMAN: Great. I think we're going to take two more questions so if you really got a zinger, then put your hand up. Otherwise you'll be later.

You sir?

MR. CURRY: My name is David Curry, and I'm with TRUSTe, and I had a question for Mike. You seem to be the only one who's done any real blocking with Sender ID, and I just wanted to know, you mentioned a statistic. Is that something that you're hard blocking now, and if so are you noticing practical issues with doing a hard block?

MR. CHADWICK: With SPF, I could recheck the message and that's where we have a lot of communication with different companies that are just -- you're testing a solution. If you still accept it and don't do anything with it and then you communicate back to the company that published the record, how do we know they're wrong or they're incorrect?

So it's kind of part of our testing cycle. We only put it out there for maybe like six or seven weeks, something like that. We're watching it. We're working with a lot of different companies, probably two a day right now, fixing their records.

So they're like, oh, we haven't even figured, and they go and fix it, and the next day their emails are coming through fine.

MR. CURRY: How soon do you think you're going to go to a bounce?

MR. CHADWICK: We are bouncing them now.

MR. CURRY: But on a test basis on a full scale.

MR. CHADWICK: It's full scale across our enterprise right now. That's why I was saying, about 18 percent of all email attached to SPF, if it's rejected, we bounce it back.

MR. CURRY: That's not what he said.

(Applause.)

MS. COLEMAN: I actually think I would like to end right there. We got applause. Thanks for having guts. That's a good close. Unless somebody has one more question, we're going to close down the shop for today. Great. Great. Thanks everyone.

(Applause.)

(Time noted: 5:15 p.m.)

C E R T I F I C A T I O N O F R E P O R T E R

DOCKET/FILE NUMBER: P044411
CASE TITLE: EMAIL SUMMIT AUTHENTICATION
HEARING DATE: NOVEMBER 9, 2004

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the tapes transcribed by me on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: NOVEMBER 24, 2003

DEBRA L. MAHEUX

C E R T I F I C A T I O N O F P R O O F R E A D E R

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

DIANE QUADE