Towards a Secure Application-Semantic Aware Policy Enforcement Architecture
(Transcript of Discussion)

Srijith K. Nair

Vrije Universiteit

Matt Blaze: How do you stop me from photographing the screen that displays my mail message?

Reply: The analogue hole is always a problem unless congress does something about it. You could play the music, put a microphone in front of the speaker, and record it; that's always going to be a problem which I don't think it's technically feasible to solve.

Bruce Christianson: You're relying on the fact that most people won't, or that the quality will be sufficiently degraded.

Reply: Yes, that the quality will be sufficiently degraded, but also the kind of system we are trying to build is basically for, like if you are employed in an organisation and all the hardware has been lent to you by that company, and you want to implement a policy on that set of machines, how do you actually go about doing it? A typical example is email.

Matt Blaze: But the email example is different from the music copying example, even though in principle it's digital rights management, because the scale of an email message is vastly different from the scale of all of the nice sounding notes of a piece of music.

Reply: The problem with email is that you can actually have a policy which says, don't forward it to ABC, but you can forward to somebody else; DRM is usually binary, you play or don't play, so this could be a bit more than binary decision-making.

Matt Blaze: Yes, I mean, what I want to do is tag the message with, don't use this message in any way that I will regret.

Reply: Yes, exactly, whistleblower.

Bruce Christianson: There's a fine line because if you're not allowed to forward it, but you can reply to it, and you can circulate the reply and the original message is included in the reply. . .

Reply: That's true, so you have to make a search of how much of the data is actually in the new message. I mean, you could always do a reply, and then change the "to" address, so the thing is, if you have a base policy enforcer,

B. Christianson et al. (Eds.): Security Protocols 2006, LNCS 5087, pp. 32–39, 2009.

you can never catch it, because at that level it's just binary data, but if it's at application policy enforcer level you can actually tell the difference between whether the "to" has been changed or not.

Bruce Christianson: So the relation between the plugins in the next layer down is, you've got to think quite carefully about it haven't you. OK. That's why you have those two extra calls.

Reply: Yes, that's right.

Mike Bond: I think if I had one high level concern it would be trying to add to a policy enforcement architecture some access control which decides who gets the benefit of policy enforcement. If you have policy enforcement for everybody, then you're going to find that your computer fills up with different applications with different policies, which either conflict with each other, or conflict with your own wishes. To give a simple example, you have a text editor which installs itself and decides to show adverts, but says that only I am allowed to read text files, and then you're stuck with that, because it's enforced its policy on your files.

Reply: Yes but if you looked at the requirements which I had before, one of the things which I wanted to have is a policy that was actually fair to all involved parties. So what happens is, if you want to open a text file, which can be opened only at the text editor, you should be told upfront that this is what's going to happen; if you want to read a mail which says "don't forward", you should be told, OK, you're going to read a mail which is "don't forward" tagged, so do you want to read it or not.

Mike Bond: Yes, what I'm saying is that as you bring in more and more units to these systems, there are decisions that you may not easily be able to go back on. OK, in the case where you've got a steady application basis, a corporation, then I guess it's not a problem.

Reply: You mean the choice of applications, like I could be opening a file in A application, or a B application, and if A sends me a lot of ads, I can switch over to B?

Mike Bond: What I'm saying is, if both A and B have a policy saying, I'm the only person who can read this sort of file.

Reply: Oh, we still want that decision to be able to be made by the user. There is still a fine line between what a user should or should not be allowed to do, I would say the user should be able to choose what application can open a file. So if I'm fed up with two applications, you can get somebody else to write an application for you, or somebody else will feel, OK, this is a good opportunity for me to write that application.

Tuomas Aura: But then you can write your own application that prints the file.

Reply: Exactly, and that's exactly the point we are trying to prevent, because your application is not trusted at any point, so if you can actually catch it at its system level calls, irrespective of what application it is, you should be able to prevent that. That's the biggest problem we are facing.

Bruce Christianson: Your argument is that you're enforcing these semantics, not by controlling which applications run, but by controlling what they do?

Reply: Yes.

Tage Stabell-Kulø: To me enforcement and fair is a contradiction in terms, because enforcement means that you decide something over me, while fair means that I decide what I want to do.

Reply: Yes, but it also brings in the question, who owns the data, which is a big question because if I send an email to you, who does the email belong to?

Tage Stabell-Kulø: It belongs to me.

Reply: But I sent it.

Tage Stabell-Kulø: Yes, you're right. But, if you give something to me, it becomes mine, I'm sorry, this is the way it is.

Bruce Christianson: Your statement is very useful because it's highly controversial. [Laughter] The idea that because I send you something that's copyrighted, I'm relinquishing all my rights over it, is perhaps naïve.

Bruno Crispo: Suppose that you have this service provider who collects personal data, it's your data, but they claim that they don't forward to any third party. In the commercial world, I give you my credit card details, the credit card is not yours, you just use it, but it's still mine, so it's not clear what is the ownership of an email.

Reply: So if I have a mobile agent and I send it to you, that mobile agent is yours, to do what you like.

Bruno Crispo: What belongs to you is the space that it takes, not the information it contains. I think as a general rule, when computer scientists act as lawyers, run for cover. The difference between the copyright of some information, and information itself, is totally a different issue, and now we're hearing lawyer talk, so we should run for cover. Copyright is a legal construct. But if I ask you in advance, look I am sending you something that belongs to me, then you can agree to give some of your hardware, some of your storage or not, accepting that actually you're hosting some piece of information that is not yours.

Tage Stabell-Kulø: So you send me an email, it lands on my desk, and you say that you can ensure that it still belongs to you?

Bruno Crispo: Well for the mail, not.

Kenny Paterson: Let me give you an example, Google mail is very popular, I set-up Google mail, Google has the right to trawl through all of my email, and through certain advertising on the website to try to sell me services, I'm very happy because I get a very large storage area, but I would rather they didn't use that email then to then market other services to me. At the moment that's all done through my trust in the Google brand and wouldn't it be nice if we could have a system which guarantees that.

Tage Stabell-Kulø: Can we trust Google to be careful with our data?

Reply: No, you can't, because there was a recent case where somebody was told to hand over all the data of his mail, even the deleted mail.

Bruno Crispo: Because there is also the owner, so even if you don't want to then you may be forced to hand it over.

Bruce Christianson: Sure, but suppose you are providing the service, and you wish to protect yourself against malicious accusations that you're forwarding mails that are not the mail that was sent to you, and so it's very useful to be able to defend yourself by saying, well I cannot do that.

Tage Stabell-Kulø: Yes, and here is the crux of the issue. Do you know of any other way of protecting information that you do not physically control other than encryption?

Tuomas Aura: Well the email's on your hard drive, but if I sent an email, it is mine. Your computers have access to it, you can read it, but you are not allowed to publish it in a newspaper.

Tyler Moore: But if I send something to you, and it's a Google email account, then Google serves as a third party which also gets this access to communication, that is the part that you want to regulate, what exactly they can do. So we're dealing with different users, we're storing this information, so you have different policies.

Tage Stabell-Kulø: Yes, at the moment I'm getting to feel that legal contracts and copyrights get things much more murky.

Chris Mitchell: It seems to me this whole conversation is invalidated by something that you said earlier on, that you are aiming at a corporate environment where the hardware doesn't belong to me even though, you mentioned, it's on my desk. In which case, this conversation doesn't apply. I suspect this is where all the money is, all these solutions are not for the home user, because we're not going to be running a policy enforcement layer on our own PCs, well certainly I don't suppose I will be. So, I think there's a great danger in getting very worked up about what the corporates do, and trying to apply it to the home user, I think this is where a lot of heat has been generated without necessarily a lot of light.

Mike Bond: I would counter-argue that corporate users have already got perfectly adequate systems, I think IRM[1] for instance, with Microsoft Office worked great in corporate environments, do they really need more?

[1] See `http://office.microsoft.com/en-us/help/HA101003661033.aspx`

Chris Mitchell: Maybe that's true, but I think they do have already difficulty in managing security because things are so much more distributed. This notebook might be part of the corporate environment, they may have bought it, and they may have installed things on it, but it's not within their firewall anymore, they've lost the little control of what goes in and out of this machine.

Tuomas Aura: Can I ask just a purely technical question [Laughter]. If you have communications like, do not print, or, do not forward, on the data, do not email, and you say you allow any application to process it. Now cutting and pasting data from one document to another, that should infect the other document with the same access controls.

Reply: Well, it's a problem of tainting. Your X-Window environment should know the policy differences, and there are ordinary environments which actually take care of it. You have a high priority window which cannot cut and paste in the low priority window.

Tuomas Aura: These windows are now applications that are trusted, but you say, I'm going to write my own applications. If I write my own applications where I can read emails, I can now take, and I can maybe draw paintings, and write books, and so on, so I accidentally cut and paste from one email that says "do not forward, do not print, do not do anything", I cut and paste a section to the book I was just writing.

Reply: Now the book is tainted, yes, but the question is, does the application actually provide the whole X-Window environment, because there is a difference between an application, and what is rendering the data onto your screen.

Tuomas Aura: I can write my own applications, I can write an application that can process any kind of data, it doesn't need to provide any windows, it just has the capability of reading one file that has access controls on it, that I receive by email, and writing another file that I created myself. Now if I can cut and paste anything from the file, with a policy on it, to this document that I have created myself, will my document be tainted?

Reply: It will be tainted.

Tuomas Aura: So do you see what happens, you will have a high watermark system where suddenly, I can never publish this book because I can't print it anymore. Because I accidentally cut and paste something from email to that file.

Reply: Yes, but that accident is always a point where you can actually leak information. Do you want to err on the safe side, or on the sloppy side?

Tuomas Aura: But that's the way we use computer systems, we process data with different tools, and we combine them, and we use pieces of data from here and there, cut and paste is one way of doing that, attaching objects is another way. If there is a policy "do not print", that means anything that is tainted with this data will have a "do not print" on it, eventually everything on my hard disk will have "do not print" on it. Unless I'm very careful not to taint other documents.

Bruno Crispo: But you cannot transfer from one application because it doesn't succeed to transfer the data.

Bruce Christianson: The purpose of process isolation is to stop you from doing that.

Tuomas Aura: But, so you're saying that there's one process for each document, and you can never pool information between two processes. But both processes are running applications that I have written, so they both cooperate as much as possible.

Reply: Yes, but if you're sandboxed, you can do process isolation independent of the processes.

Tuomas Aura: So this means there is no free communication between any two processes.

Bruno Crispo: Yes, because otherwise you break the policy.

Bruce Christianson: Any IPC must go through a system call that is subject to checking against the policy.

Reply: This argument is, the policy doesn't say "don't copy", the policy says "don't print", so in that case, you're right, the policy taints the next file, but the problem is, do you want to allow that, or not allow that.

Tuomas Aura: It's not a question of me, it's the person who sends the email, isn't it, who sends the policy.

Bruce Christianson: Perhaps he sent you an email saying "don't publish".

Tuomas Aura: Is the user allowed to do anything to recover from this?

Michael Roe: Isn't this just a mandatory access control? You've got this classified document, and you are writing a book that's currently unclassified at a compartment node at my workstation, if you were to cut and paste a paragraph from the classified document you can do two things, you can either upgrade your book so it becomes a classified book, or you can not cut and paste, and the workstation enforces that.

Bruno Crispo: Yes, but it's not mandatory in the sense that here every user has his own policy for his own documents. There is not a system that actually can force a mandatory access control here, you have a policy attached to a document, it's not exactly the same. But it is true that this goes really against what people are used to nowadays because, for example, you have a piece of a Word file, and you can go from Powerpoint to Entourage to Office, essentially there is no compartmentalisation of any software applications.

Tuomas Aura: But in these kind of emails you won't have just one compartment, you have as many as there are emails in your inbox.

Bruce Christianson: Yes, yes.

Tuomas Aura: And this makes things really complex, because there are…

Reply: No it doesn't.

Bruce Christianson: No, it's only since the complete lack of self control that arrived with personal computing that we've had this idea that processes are expensive and managing a lot of them is unnecessarily difficult. This is more of a return to the way things worked in the mainframe world.

Tuomas Aura: But you're managing many different security policies.

Bruce Christianson: Where this gets really interesting is exactly where you're composing policies. For example, if I have a project team that's working for a company that are sub-contractors to another company on this project, but the company's bidding against them on a different project, and I want to be sure that they're complying with the security policy for this project, and they're not under management pressure going to, for example, cut corners on testing, or do something that would make us look bad in the other contract that we're bidding for. Here they've got to comply both with their own company's policies, and with the policy of the project in terms of the tender we bid for, now that seems to me to be something where an architecture like this does have something to offer, and it's precisely because you are composing policies. You can't do that in a flat environment.

Tuomas Aura: So that means every time I cut and paste, or any time I move information from one document process to another, I need to read through this complex security policy to check what kind of restrictions there are in place?

Bruce Christianson: You don't have to do it yourself, there's software that does it for you.

Tuomas Aura: Sure, software does it, then I accidentally…

Bruce Christianson: Well if you can do it accidentally, your software's not as friendly as you thought.

Matt Johnson: Can I ask you where you're envisaging all these processes running, is this on a Unix style system where it's all running on a remote server, or are you talking about also having it running on the local machine?

Reply: It's actually much easier to enforce on a local machine.

Matt Johnson: Well can you explain that, because if you've got a local process that's got access to a file, and has all the necessary information to do whatever decryption it needs to get at the data, then why can't I write my own code, possibly by reverse engineering the processes running on my local machine, and get at the data. It seems to be the same problem with encrypted DVDs, that if everything's got to know the key to ever get at the data, then sooner or later somebody's going to leak the key.

Bruno Crispo: Be careful, because the assumption here is the hardware is trusted and the operating system is trusted, so it's limited.

Reply: You see the thing is, first of all you find that you can write your own application, and that you can have access, if you are a user yourself, and you can allow a particular application to have access to the data, the question is, what can you do with the data, it's there in your memory, but what else can you do with it.

Matt Johnson: OK, so this could be some sort of tamper proof computer, so that I can't look at things sector by sector.

Reply: Yes, why not, I mean, it's already there, because when you do process isolation you have all this hardware base support already available to you.

Bruno Crispo: You need also a trusted operating system, because otherwise you're in trouble as well.

Matt Johnson: Well do you need a trusted operating system, so that you can't just rip the hardware out and go and take it somewhere else.

Reply: No. The actual implementation could have things to do with decryption, so even if you take all the hard disk to another machine, all you will have is encrypted data. The actual technical implementation is still being worked on, so how you actually enforce it, is still a problem.

Matt Johnson: But if my machine needs to be able to decrypt what's on the hard drive when it boots up, then my machine needs to know what the key is, and if my machine knows what the key is, then I can get that key out of my machine.

Reply: An example would be TPM. I mean, a TPM has a hardware-based key which you can't get from outside.

Matt Johnson: OK, so you need tamper resistant hardware?

Reply: Yes. You basically build on top of hardware which is processed.

Matt Johnson: You can never email one of these documents to somebody who's using a normal PC at the moment.

Bruno Crispo: Again one of the assumptions is that we don't care about backwards compatability, yes, that's true. Otherwise, it's no way. The world is already difficult enough as it is.