

Global ID, Trusted Systems, and Communications Markets

Jonathan Weinberg *

This paper examines the implications of technology identifying each consumer by a single globally unique identifier, or GUID. Such technology allows Internet content providers straightforwardly to identify the consumer originating any given packet stream, and to correlate incoming payment (and other) information to the information and entertainment that the content provider releases to that consumer. These systems promise to give content providers sharply expanded powers to discriminate among consumers. The paper suggests that, on balance, this result would be a Bad Thing. Further, the result could be avoided if content providers relied on more sophisticated techniques to manage access to their information goods.

I. Rights Management and Trusted Systems

The most important force today driving information flow from consumers to content providers is rights management. The term “rights management” is commonly associated with the protection of intellectual property rights, but it need not be so limited: One can think of it as covering any technological means of controlling public access to, and manipulation of, digital resources. That sort of control is basic to any system of networked computing. At the heart of Unix, for example, is the concept of permissions, which define *which* users on a network can take *what* actions (read, write, execute) on *which* files and directories. Networking wouldn’t be very practical unless there were a way of defining (and limiting) the set of people who can have access to particular documents and other network resources. Rights management techniques in that sense are simply a form of network security.

Those techniques demand a reliable way to match usernames with real-world individuals. After all, it’s the individual, not the username, whose access to files is at issue. In the old days, when mainframe computers ruled the world, system administrators had little difficulty associating the individuals using their systems with unique usernames, and thus using permissions or similar file access rights to enforce that aspect of system security. The systems administrators had themselves assigned those usernames to the individuals in question. The situation wasn’t much different for a self-contained local area network.

But the Internet changed things. There is no reliable automated way, under current technology, to tell what individual is associated with any given username on an Internet-connected network. In-

deed, even your own computer does not know who you are; if you tell your PC that you are Napoleon or Joan of Arc, it has no reason to disbelieve you. The ordinary Internet architecture, thus, stymies attempts at rights management beyond a given network. It provides no convenient set of options in the middle ground between blocking access by anyone outside one's own network, and granting access to everyone in the world.

How can one extend sophisticated file access rights beyond the controlled network environment into the Internet universe at large? Put another way, how can a local server extend secure control over the many interconnected networks that make up the Internet? To do that, it must be able reliably to identify everybody out there seeking access to its files (or, at least, everybody to whom it is willing to grant access), and then be able to sort those persons by whichever of their characteristics it deems relevant. That is to say, it must have some way of reliably associating incoming packet streams with identified real-world individuals, and it must have — or be able to collect — enough information about each of those individuals in order to be able to implement a set of rules determining whether to grant access.

One way for a content provider to accomplish these tasks is to allow access only if the recipient's computer (or other device) can prove, through cryptographic techniques, that it incorporates hardware and software, meeting security specifications approved by the content provider, that will enforce rules specifying the terms under which individuals can access, and use, particular digital content. In such a case, technologists refer to the server and recipient as being part of a *trusted system*. The server can rely on "trusted" elements of the recipients' device to identify the recipient reliably, to transmit only accurate information about the recipient, and to limit the recipient's ability to manipulate any content it receives from the server in ways that exceed its authorization.

Trusted systems enable the sophisticated network security I discussed above, because they give the content provider a way to verify the authenticity of any message it receives claiming authorization to read a digital work. But they give content providers broader capabilities. They allow the content provider to make the works available only to persons the content provider knows to have paid for access. They allow the content provider to prevent the recipient from passing usable copies of the works to unauthorized persons (say, those who have not paid). They allow the content provider great flexibility in specifying the nature of the event that will trigger a payment obligation — for example, a content provider could allow a consumer to download a work for free, but to pay each time she reads or listens to it. In short, trusted systems have the capability to be an extraordinarily effective (and profitable) means of controlling, and rationing, access to works of information and entertainment.

In order to implement large-scale trusted systems, however, the computer industry must develop technology to feed reliable identifying information about consumers back to content providers. The next section of this paper describes one such technology.

II. Intel and the Processor Serial Number

Early in 1999, Intel — which manufactures the vast majority of the chips powering personal computers today — introduced a technology that it described as the foundation for a whole new world of trusted systems: the Processor Serial Number, or PSN. The PSN is a unique identification number burned into each CPU as part of the normal manufacturing process. Intel announced plans to incorporate the PSN in all of its products, including not only its Pentium III chips for personal computers, but also the microprocessors embedded in devices such as televisions, stereos, telephones and "Internet appliances." Applications running on any device equipped with a PSN can read the unique identification number and transmit it to any requesting remote server. Such a system could provide the foundation for reliable flow of identification information from every consumer to Internet-based content providers.

In introducing the PSN, Intel vice president Patrick Gelsinger explained that the company was shifting its vision from that of “a world of a billion connected computers” to that of “a billion *trusted* computers.” A vision of a world fully populated with a myriad of personal computers, each communicating with the rest, he explained, is insufficient unless those connected computers are trusted, and the first step on “the road to . . . trusted connected PCs” is the PSN. Because each computer’s PSN is unique, he continued, the PSN provides a hardware framework for treating the home PC as part of a trusted system — that is, to allow servers on distant networks to authenticate the identity of a home PC user, and administer authenticated permissioning and rights management. It could thus create a “trusted virtual world” for secure virtual enterprises, business-to-consumer electronic commerce, and secure delivery of high-value digital media content (movies and audio).

Gelsinger explained that the PSN, “enabl[ing] platforms and the users that are on those platforms to be better identified,” was Intel’s first building block in constructing this system. “You think about this maybe as a chat room, where unless you’re able to deliver the processor serial number, you’re not able to enter that protected chat room[; it provides] a level of access control.” Gelsinger announced plans to add significantly to the capabilities of the PSN and the Common Data Security Architecture the following year, “allowing . . . trusted access, adding authenticated permissioning to PCs, [and] increasing levels of capability” in the security architecture. He announced plans to add capability in 2001 relating to “platform and peripheral integrity,” thus “accomplishing the trusted transactions [and providing] a platform strong enough to bring all forms of valuable content to the PC.”

The PSN excited considerable controversy. Privacy advocates requested that the Federal Trade Commission initiate an inquiry, and followed that up with a complaint formally asking the Commission to halt distribution of the Pentium III as a violation of individual privacy. The Electronic Privacy Information Center announced a boycott of Intel. Large PC makers responded by announcing that, in shipping Pentium III machines for the consumer market, they would set the BIOS (the first software instructions a computer loads when it boots) so as to make the PSN invisible to most programs. Intel announced plans to release software patches that consumers could use to do the same thing.

Chastened by the public reaction to the PSN, Intel retreated and regrouped. It has not mentioned the PSN in any public statement in recent months; instead, it has sought to focus attention on its new Trusted Computing Platform Alliance (TCPA) initiative. Like the trusted-computing program Intel had earlier announced, the TCPA is seeking to deliver an “enhanced [hardware] and [operating system] based trusted computing platform” to ensure, among other things, “platform authentication” — to provide a standard way for outsiders to query a computer and establish its owner’s identity, thus establishing “confidence in interacting with [that] platform.” This time, however, TCPA statements emphasize that computer owners must control their personal information and the system’s authentication capabilities. The TCPA has not yet released to the public a specification with the details.

III. Trusted Systems and the GUID

I. Introduction

I want to explore some of the social implications of implementing, on a widespread basis, trusted systems based on identifiers such as the PSN. In that connection, it seems to me that two characteristics of the PSN are notable. First, the PSN is keyed to the holder’s *identity*, rather than his characteristics. It enforces a particular model of trust, in which, to learn the characteristics of a particular would-be information recipient, a publisher first ascertains that person’s identity and then looks up the characteristics associated with that identity. For the PSN to be used as the basis for a trusted system, the content provider must correlate the PSN with its other data relating to the individual owning that computer, by tying all of that data to the single identifier that the PSN represents. This

model stands in contrast to a more privacy-protective approach, in which a person can present credentials verifying certain characteristics (such as country of residence) without necessarily disclosing his identity at all.

The second notable characteristic of the PSN is that it is a *common* identifier. That is, it is well-suited to being used by different information collectors across a wide range of unrelated transactions, increasing the ease with which a wide range of information about a person can be aggregated into a single overall dossier. The greatest obstacle to efficient aggregation and manipulation of data today is the need to reconcile inconsistent formats and identifiers; a standard, common GUID can eliminate that obstacle. To the extent that a variety of content providers and other merchants have each collected information tied to individual PSNs, it is a simple matter to compile those files into larger databases.

2. Privacy

It is straightforward that trusted systems based on common identifiers — in which the user’s computer identifies itself during every transaction, to anybody who asks — are pernicious from a privacy perspective. They allow the user to be tracked, rather more easily and thoroughly than is possible under current technology, through cyberspace. Under such an architecture, a much greater proportion of ordinary transactions would require consumers to present unique identification numbers that will in turn be digitally linked to a much wider range of personally identifiable information.

Systems facilitating the close tracking of content — of what people read, view or listen to — seem particularly problematic. All of these are the constituents of human thought. In the analog world, sales of copies of works of information or entertainment commonly are cash-based, leaving no paper or electron trail. The copies themselves have no surveillance capabilities, and cannot report back to their makers. The copyright owner, indeed, collects no information about the user at all. Trusted systems threaten to abandon those rules, facilitating the monitoring of individual thought. They raise the specter of the Panopticon, and of subtle and not-so-subtle pressures on individuals to conform and to eschew “dangerous” works of information and entertainment.

3. Communications Policy

A more subtle set of consequences relates to the effects of this technology on the economics and politics of content markets — that is, on speech. In the Old Way of Doing Things, technical inefficiencies made it difficult to disseminate speech to a dispersed but tightly controlled group of folk. Rather, there was some leakage: If you wanted to disseminate speech, you had to give up some control over its dissemination. For example, once a content owner distributed a copy of a work, it had no technological means of preventing the owner of that copy from selling, loaning, privately displaying, or giving away that copy as he chose. And the copyright law’s “first sale doctrine” denied content holders the ability to impose such restrictions within the four corners of the copyright law. That set of limitations on content owners’ effective rights helped democratize access to content. It allowed gratis redistribution of, and secondary markets in, copies of the works.

Small-scale, decentralized reproduction has long been a fact of life in markets for information, entertainment and computer software. People copy music tapes and CDs for themselves, family and friends; they photocopy magazine articles; they allow family and friends to use, and copy, their computer software. They persist in doing so, notwithstanding the best efforts of the copyright industries to convince them that it is illegal, largely because they find it hard to believe that this is something the law does or should proscribe. And that system seems to work — at least, it hasn’t obviously injured producer incentives to create in any palpable way.

The non-trusted system world is also characterized by a lot of sharing, in the vernacular sense, that doesn't implicate copyright law. People *lend* each other analog copies of protected works, and read, watch or listen to works they have borrowed, all without implicating the copyright laws at all. At least in a static analysis, both of these sorts of sharing are good, since they increase the distribution of the informational work (and thus social benefit), without any social cost. Put another way, sharing allows distribution of the work at near the optimal demand price (marginal cost), which in this case is close to zero. The most successful institutions in American life today based on such sharing are public libraries, which were established precisely so as to enable large-scale sharing of analog works. In the trusted-system world, however, the content provider's enhanced control over access to the work would allow content providers to sharply limit both of these forms of sharing.

Another change, with more ambiguous results: The trusted-system world seems well-suited to facilitate discrimination on the part of the content provider. In particular, it would facilitate price discrimination, in which a content provider asks different consumers to pay different average prices, unrelated to the provider's own costs. This discrimination is possible for two reasons. First, trusted systems allow the seller ability to prevent (or limit) arbitrage. That is, the seller can prevent buyers from reselling the information or entertainment to someone who would otherwise be willing to pay the content provider the higher price. Second, the seller can set prices in a way that reflects individual consumers' willingness to pay. A GUID-based trusted online architecture should make it possible for a content provider to link each consumer with a wide range of personally identifiable information. Thus, when the consumer presents her PSN in order to gain access to a digital work, the content provider will be able to pull up other information associated with that PSN in order to make a judgment about the particular consumer's willingness to pay. Alternatively, the content provider can shift its payment model from the "sale" model prevalent today (in which the consumer buys a copy of the work, and can then read, listen to or watch that copy unlimited times without further payment), to a system in which the customer pays a smaller amount for each occasion in which she reads, listens to or watches the work. This allows the content provider to collect more money from those customers who want to view the work multiple times (and presumably are willing to pay more for that ability), and less from those who want, say, to view the work only once. The difference between those prices is largely unrelated to the content provider's own costs.

Some have argued that this price discrimination is a good thing. Price discrimination in information goods is socially useful, the argument runs, because it increases the distribution of the good. Without price discrimination, the content provider must charge a single market price, and people unwilling to pay that price will be shut out of the market. If the content provider can engage in price discrimination, by contrast, it can charge every consumer the price (and only that price) that she is willing to pay, thus simultaneously maximizing profits and maximizing the number of people who will be exposed to the information and entertainment in question.

The matter, though, is not nearly so straightforward. In thinking about whether the price discrimination that trusted systems would enable would be a good thing, we need to ask the question "compared to what"? One of the key reasons that trusted systems are good at enabling price discrimination is that they sharply decrease sharing; they are designed to eliminate any redistribution of the information good outside of the control of the content provider. That is, price discrimination allows the market to extend to consumers with lower willingness to pay, but at the expense of cutting off *existing* means, through sharing and secondary markets, of getting the information or entertainment at low or no cost to some of those same consumers. Indeed, sharing is from a static perspective a more efficient way of allowing the market to extend to those consumers, since it makes the good available to them at a price more nearly approaching the zero marginal cost of supplying it to them.

Secondary markets (that is, redistribution of the information good after its first sale, outside the control of the initial seller) can do the same job as price discrimination of getting information goods, at lower prices, to lower-valuation users; that's what used bookstores are all about. The price discrimination that trusted systems may make possible, thus, may not increase the number of consumers getting the good at all; it may simply ensure that the low-valuation consumers receive the good from the initial seller rather than someone else.

This point is open to a variety of counter-arguments. First, it might be argued that price discrimination will do a better job of getting the information or entertainment to low-valuation users. Many low-valuation users may not have the opportunity to gain access to the work through resale or sharing. Yet public libraries, at least, are set up precisely for the purpose of getting information works, gratis, to users unwilling to pay the price set by the market. Viewers interested in viewing a work only once (and waiting until it is available) are able to borrow from the library; those interested in viewing the work multiple times are more inclined to buy it. This is precisely the sort of result price discrimination is supposed to achieve. By contrast, it is unclear to what extent price discrimination in practice can in fact achieve the advantages theory promises for it. It is difficult to gauge consumer preferences precisely, and publishers are unlikely to drop prices too far based on guesses about a particular class of consumers' willingness to pay. DIVX's splashy failure, further, should cause us to have some doubt about the enthusiasm with which ordinary folks will embrace pay-per-view plans for digital works.

Next, one might argue that this analysis overlooks the dynamic impact of sharing, and the nature of secondary markets in digital works: Sharing and resale don't generate revenues to the content provider, so they don't provide incentives and don't stimulate production. More baldly, one might argue that my discussion is in essence an argument for piracy — which will certainly lower prices to the consumer, but at the cost of diminishing incentives to produce. Secondary markets in the digital world, the argument runs, may involve large numbers of illegal perfect copies. Sale of those copies cuts directly into the profits, and thus the incentives, of the initial producer.

I don't contest that producer incentives are necessary. Publishers have to be able to sell information goods at a price sufficiently above marginal cost, for sufficiently long, to enable them to recover their fixed (first-copy) costs. Otherwise, they'd lose money. To that end, there have to be sufficient entry barriers limiting other folks' ability to sell those works as cheaply. We don't know, though, how much in the way of incentives producers need. While increasing producers' ability to extract rents from their information goods increases their expected profits and thus, perhaps, their investment, it also increases the deadweight loss on society that the inefficiently high prices impose. If publishers have adequate incentives even without the extra rents that price discrimination gives them, then we *may* get a better social result by reaching lower-valuation users through secondary markets, sharing, or even (some degree of) piracy than through the increased control that trusted systems bring.

There's a connection between the power of information providers to identify consumers, and to thus to discriminate, and media concentration. To the extent that sellers' ability to price discriminate will rest on their access to personally identifiable information about buyers, publishers with access to those databases will have a competitive advantage over those who do not. This may have two negative effects. First, it will tend towards concentration in media markets — and, to the extent those markets are characterized by winner-take-all dynamics, will help determine who those winners are. Second, it will increase the value of the dossiers, and thus increase the commercial pressure on privacy.

The control facilitated by GUID-based trusted systems may allow other sorts of discrimination as well. Most generally, it will enable producers more nearly to pick and choose who will be allowed to view, or read, particular works. Given the power of a common identifier such as the PSN to facilitate the association of a wide range of information with a given personal identifier, producers could in

theory use these tools to allow access to a speech work only by persons who live in preferred zip codes, or have certain levels of family income, or are white. There may be only limited circumstances in which a mass marketer of entertainment and information would have incentive to do so: most obviously, perhaps, on the basis of ideological motivations, or if particular content gained cachet from only limited distribution. It seems disturbing, though, from a free-speech and communications policy standpoint, to see extensive social investment in a technology built around the ability to *prevent* the movement of speech and information to the public at large.

For the most part, today, content producers and consumers share control over the uses and dissemination of speech works. Content producers have extensive control by virtue of their ability to produce, and license, the technological artifacts (such as film reels) embodying those works, reinforced by the rights granted them by the copyright law. Consumers have some control as well, by virtue of their own abilities to use, copy and manipulate such works in ways that the copyright law either does not forbid or expressly privileges, or in ways that have been effectively immune from copyright enforcement. And because these are speech works, that distribution of control has political consequences. It shapes the overall movement of information and expression within society. The rise of GUIDs and global trusted systems threatens to shift that control.

IV. Identification and Credentials

GUID-based trusted systems, in short, seem to have a variety of undesirable social consequences. And yet, one might think, they are unavoidable if we are to allow content providers control over exploitation of their works in the networked digital environment. That statement, though, is not correct. In fact, the Internet's architecture can support trusted systems — and concomitant control by content providers over works of information and entertainment — without any need for GUIDs or other common identifiers.

Recall the original concern driving industry plans for unique identification of Internet-connected computers/consumers: content providers wish to be sure that a packet stream requesting access comes from a person with characteristics that entitle him to access (say, that he has paid). One way to accomplish that result is to tag every computer/consumer with a single identifier that both shows up in the packet stream requesting access, and allows reference to a database of consumer's characteristics (say, whether they have paid). But that approach conveys much more information to the content provider than the content provider actually needs.

The content provider needs some way to verify that the user has specific credentials (that the user has paid, say), or that he has some other characteristic that the content provider desires in its readers. Establishing the user's identity is an instrumental step towards verifying his credentials. Yet it's well-established in the cryptography literature that one can prove credentials without proving identity: that's the basis for anonymous digital cash. A person, for example, can interact with other entities through a *pseudonym* — a name that is reliably associated with that individual in a particular context through cryptographic techniques, but cannot be associated with other names the person uses in other contexts. The word "pseudonym" sounds vaguely disreputable, but the goal is simple and (usually) honorable: It is to allow the user to enter into transactions and relationships in which they can be held accountable, without creating the opportunity for data miners to collect the universe of transactions they enter into a single global profile. It is consistent with current rights-management technology to build structures under which content owners interact with consumers anonymously or pseudonymously, without sacrificing content owners' ability to enforce contractual restrictions.

Such systems would address the privacy issues raised earlier in this paper, by making impossible the aggregation of a user's information across unrelated transactions. They would not make it impossible for a content provider to discriminate among users, but they would make that process more

open and public. Because the content provider would not know any information about the user that the user did not provide to it, it could discriminate on the basis of a particular characteristic only after expressly asking the user to provide credentials relating to that characteristic. Content providers would be reluctant to seek information where such requests would be unpopular in the marketplace or the forum of public opinion.

I do not mean to suggest that systems protecting user privacy would be the first choice of content providers. For the reasons set out earlier in this paper, content providers may find such systems significantly less profitable, and hence less desirable, than those that give them access to a greater range of user information. From the perspective of social policy, on the other hand, building trusted systems around GUIDs is both unnecessary and undesirable.

* Professor of Law, Wayne State University. I owe thanks to Phil Agre, Karl Auerbach, Lorrie Cranor, Jessica Litman, Neil Netanel and Joel Reidenberg. This paper is a shorter (and footnote-free) version of *Hardware-Based ID, Rights Management, and Trusted Systems*, forthcoming in the Stanford Law Review. The most recent version of the long paper is available at <<http://www.law.wayne.edu/weinberg>>. Earlier versions were presented at a Haifa University Conference on the Commodification of Information and at the Telecommunications Policy Research Conference.