Perils and Pitfalls of Practical CyberCommerce

The Lessons of First Virtual's First Year

to be presented at:
Frontiers in Electronic Commerce
Austin, Texas
October, 1995

Nathaniel S. Borenstein <nsb@fv.com>
John Ferguson <ferg@fv.com>
Jerry Hall <gwhiz@fv.com>
Carlyn Lowery <lowery@fv.com>
Rich Mintz <mintz@fv.com>
Darren New <dnew@fv.com>
Beverly Parenti <beverly@fv.com>
Marshall Rose <mrose@fv.com>
Einar Stefferud <stef@fv.com>
Lee Stein <lstein@fv.com>
Carey Storm <cstorm@fv.com>
Ed Vielmetti <emv@fv.com>
Marc Weiser <maw@fv.com>
Pierre-R. Wolff <pierre@fv.com>

I. Introduction

Unlike many would-be players in the field of Internet commerce, First Virtual (tm)[1] chose to announce its payment system only after it was fully operational, and to operate it initially with relatively little publicity hype, while learning from the experience of its use. In its first year of operation, it has experienced exponential growth, and the company has gained substantial experience with and insight into the nature of Internet Commerce. In this paper, the First Virtual team discusses the lessons we have learned from a year's experience with the actual operation of an Internet commerce system, and the prospects for the future.

This paper begins with a short description of First Virtual and its Internet Payment System, which may be skipped by those already familiar with it at the conceptual level. Next, we consider the lessons learned, focusing on five key areas: the organizational aspects of an Internet service company, the need for an Internet-based intermediary in the payment process, the security and administrative issues involved in operating an Internet commerce server, the customer service issues in dealing with a user community as diverse as the Internet, and, finally, the myths and realities surrounding the use of cryptographic technology for Internet commerce. Finally, we look to the future, with

---

[1] First Virtual, Virtual PIN, and InfoHaus are registered trademarks of First Virtual Holdings Incorporated.

projections about the future evolution of First Virtual's system in particular and Internet commerce in general.

II. What is First Virtual?

First Virtual Holdings is a company that was formed in early 1994 to facilitate Internet commerce. The first product offering from First Virtual was an Internet payment system, which was developed quietly and publicly announced as a fully-operational open Internet service on October 15, 1994.

First Virtual's system differs in many ways from all other proposed approaches to Internet commerce, most notably in the fact that it does not rely on encryption or any other form of cryptography to ensure the safety of its commercial transactions. Instead, safety is ensured by enforcing a dichotomy between non-sensitive information (which may travel over the Internet) and sensitive information (which never does), and by a buyer feedback mechanism built atop existing protocols.

In a nutshell, First Virtual's payment system is built on top of pre-existing Internet protocols, notably the SMTP/RFC822/MIME (email), telnet, finger, FTP (file transfer) and HTTP (Web) protocols. Because those protocols are "insecure" in the sense that they carry no strong proofs of identity, it is necessary to design a payment system in such a way as to provide much stronger guarantees. While others have focused on achieving this goal using cryptography, First Virtual designed a higher-level protocol based on email call-backs.

In the First Virtual system, a buyer and seller may meet and decide to transact business in any manner they desire -- while this often occurs when a buyer browses a seller's Web page, it also frequently happens by email, FTP, Internet Relay Chat, or even off-net entirely, and it could easily happen in the future via protocols that do not exist today. Once the buyer and the seller have an intent to do business, they submit a transaction to First Virtual. That transaction can be submitted via standard email or via a new protocol, SMXP, designed by First Virtual for real-time exchange of MIME (email) objects.

When First Virtual is asked to process a financial transaction, it looks up the buyer's Virtual PIN (account identifier) in its database, and finds the buyer's electronic mail address of record. An email message is dispatched to the buyer, asking the buyer to confirm the validity of the transaction and his commitment to pay, which the buyer can respond to with a simple answer of "yes", "no", or "fraud". Only when the buyer says "yes" is a real-world financial transaction actually initiated. Simple attacks based on Internet "sniffing" are rendered unappealing because their value is sharply limited by the fact that a Virtual PIN (tm), or First Virtual ID, is not useful off the net, and require email confirmation for use on the net. More sophisticated attacks require criminals to break into the victim's computer account and monitor the victim's incoming mail, a crime that is much more easily traced. It is also worth noting that such a break-in would also probably yield access to the victim's encryption keys in any commerce schemes that make use of public key cryptography for encryption.

In First Virtual's system, the valuable financial tokens that underlie commerce -- notably credit card numbers and bank account information -- never appear on the Internet at all. Instead, they are linked to the buyer's Virtual PIN by First Virtual when the customer applies for a First Virtual account, a procedure that involves an off-Internet step for the most sensitive information. Currently, the sensitive information is provided by either an automated telephone call (for buyers to provide their credit card number) or by postal mail (for sellers to provide their bank account information). However, it would also be possible to provide the Virtual PINs automatically en masse to buyers, e.g. by direct mailing from the credit card issuers as is done with traditional ATM PINs.

The exclusion of the most valuable (to criminals) information from the Internet data stream eliminates any need for encryption, which in turn eliminates the need for any non-standard software on the buyer's end. Ordinary email -- which effectively represents the lowest common denominator of Internet connectivity -- is all that anyone needs in order to participate. The simplicity of this approach gained First Virtual more than a year's head start in the marketplace over the encryption-based approaches, and greatly lowered the entry barrier to anyone wishing to become a First Virtual user.

Another unusual feature of the First Virtual system is that it is explicitly designed for entrepreneurs. There is no screening process for sellers, allowing anyone on the Internet to open a new business. The system even includes an automated information server, the InfoHaus (tm), that will (for an additional fee) make information continuously available for sale by Web, FTP, and email, even for sellers who do not have their own Internet servers.

Full details about the First Virtual system are available elsewhere, as documented in the bibliography. In this paper, we will concentrate on the lessons we have learned from operating that system, which has been processing real money transactions, for an initial year. However, the system is sufficiently different from most other proposed approaches to Internet commerce that we have included a rather lengthy list of commonly-raised concerns about First Virtual, and our responses, as Appendix A to this paper.

III. What Have We Learned?

A. Organizational Issues

First Virtual has attracted some notice as an extreme example of a "virtual company". The company was certainly unusual in its initial organization: The four founders lived in San Diego, Orange County, Silicon Valley, and northern New Jersey. We promptly hired additional team members in distant parts of the same and other states. There were no physical offices until 15 months after the company was founded (8 months after the system became operational). The servers were set up in a high-security EDS machine room in a suburb of Cleveland; the data 800 number was answered in Atlanta, Georgia; the voice 800 numbers started out in Portland, Oregon, but were then changed to move around from city to city. Marketing was handled from Washington, D.C., and public relations from San Diego. The company hired lawyers in San Diego, Los Angeles, Chicago, New York, Washington, and Cheyenne. Legally, First Virtual is a Wyoming

corporation. This is, to say the least, not a typical corporate organization!

Some aspects of this decentralization worked well, and were quite fun. Certainly it was always fun to tell the story of our "virtual office," as in the previous paragraph! But there were serious problems as well. While three of the four founders were long-time Internet veterans, one was not, and approximately half of the early employees (all the non-technical ones) were Internet "newbies" who had to learn the ropes of working with others completely via the Internet. This is a non-trivial endeavor. The larger the company grew, the more seriously its productivity was impeded by communications difficulties, which ultimately led to the decision to consolidate the bulk of operations -- and particularly new hires -- in a small number of offices.

The biggest problems in running a distributed company were the more "mundane" aspects of any corporation -- administrative tasks, scheduling meetings, making presentations to customers, and so on. There were a frightening number of "near misses" in which people were told of important meetings or discussions at the last minute, and an appalling number of emergency red-eye flights. It was much harder to gather people together for informal brainstorming sessions and other creative gatherings. The distributed nature of the company made it difficult to ensure that the company would speak with a unified voice in its public statements, and to avoid wasteful duplication of efforts. It is also far harder to integrate new hires into a virtual environment, particularly if they are not by temperament the kind of independent workers who work best in such an environment.

More specifically, the actual supervision of remotely-located employees was a constant management challenge. The more distant these employees were from the initial founding and vision of the company, and the less clearly they understood the "big picture" of the company's strategy, the less likely they were to be able to execute their jobs productively without close supervision. This, in turn, was reflected directly in the degree to which their remote location was perceived as an impediment to their productivity.

Given these problems, it is tempting to say that "virtual companies don't work." This is an oversimplification, and an irrelevant one in any event. First Virtual, in particular, could not have been created any other way. Its four founders were well-known, extremely successful people who lived in four different parts of the country, and it was never a serious possibility that three of them would relocate in order to start a highly-speculative new venture. (Later, as the company grew, some such moves did take place, which leads to the speculative generalization that a distributed organization might be a reasonable starting point in a larger corporate evolution.)

More generally, the fact is that almost any Internet service company will by nature be somewhat "virtual", if only because of the need to support fully international operations. If you're going to be able to communicate with Internet-based customers around the world, in many languages, it is almost inevitable that you will end up with operations spread out to many countries, connected to each other primarily via the Internet. Thus the right question to ask is not "should an Internet company be virtual?" but rather "How virtual should an Internet company be?" or perhaps "How can the advantages of a

distributed company be maximized and the disadvantages minimized?"

What worked best in our virtual company were creative projects executed by small, strongly-motivated, highly-skilled teams. The basic technologies in First Virtual were all created by such teams, whose members never shared an office. However, the need for communication and clear task delegation among the team members argued for regular in-person meetings. Two-day monthly staff meetings, scheduled on a totally regular basis for the same days each month, have proven sufficient for such tasks.

Another ultimate strength of our operation, despite occasional problems, was the customer support system. Because all of First Virtual's customers have electronic mail, First Virtual is able to do nearly all of its customer support over the Internet. Our customer support operators are distributed across the United States, but this has not proven to be a problem. In general, the operators have worked well, and customer service has functioned well without paying rent for any office space.

One human and social benefit of a virtual company with distributed customer support is that it creates a set of jobs requiring a high level of mental skills, but which can be performed by people with severe physical disabilities. For example, First Virtual's senior customer support representative, one of the authors of this paper, is severely disabled in a manner that might inhibit his employment in many traditional work environments. By computer, from his home, he communicates using voice dictation software, and has interacted with thousands of First Virtual customers who never had any inkling that he was disabled at all. We believe that, just for this benefit alone, it is well worth tolerating some of the more challenging aspects of a virtual corporation.

As the customer support staff grew, however, it became clear that while skilled customer service operators work well remotely, training is made more difficult by distance. Accordingly, a major current focus of the customer service department is the production of improved training materials for new operators.

An intangible factor that requires special attention in a virtual environment is employee morale. It is relatively easy for an employee working remotely to come to feel "out of touch" with the company as a whole. Regular meetings are helpful in this regard, as are frequent phone conversations. (All senior management employees were required to get 3-way calling service, and they often chained together several 3-way calls as an inexpensive mechanism to establish larger conference calls.) The customer service department is also contemplating morale-boosting incentives (e.g. a "silly question of the week" contest) that will facilitate friendly competition and communication among the customer service operators, whose entire job consists of dealing with the system's "rough spots".

In short, having everybody together at a single site is absolutely not a prerequisite for doing business on the Internet, which should be a relief to anyone contemplating serious international operations. However, a distributed operation carries some very specific pitfalls in terms of communication, efficiency, and motivation, which need to be understood and addressed by management early on. It also seems very compelling to try

to centralize those operations that can be centralized, such as marketing, operations, and corporate administration.

B. The Need for an Internet Intermediary

One complaint that has been voiced about both First Virtual's system and several other proposed approaches to Internet commerce is that they create a new intermediary between the customer, the merchant, and the financial institutions. Our experience to date strongly suggests that this is not a bug, it is a feature, and that all parties involved will increasingly see the necessity of such an intermediary as the nature of Internet commerce becomes clearer.

The simple fact is that the Internet is a complex set of technologies and services that simultaneously make commerce possible and also form a barrier to the conduct of that commerce. The distributed, anarchic nature of the Internet makes certain classes of service oddities inevitable, including temporary partial network outages, total or partial communication failures either unidirectionally or bidirectionally, subtle incompatibilities between software on the buyer and seller end, and much more.

What is often overlooked is that from the buyer's perspective, the following two situations are indistinguishable:

-- A technical failure, possibly even one caused by an invisible intermediate third party, that prevents a reputable merchant from either delivering paid-for merchandise or notifying the buyer of its non-delivery and the refund procedures.

-- An unscrupulous merchant who defrauds his customers for a quick profit.

In our experience, the first case is far more common, but buyers are remarkably quick to assume the second case. This is in part human nature, and in part due to the strangeness of cyberspace business relationships, in which one sends money to some unseen person on the other side of the planet.[2]

Customers naturally expect and demand that the provider of payment services will mediate such situations and help to resolve them. Whoever performs that service is, ipso facto, a new intermediary in the payment process, to facilitate the resolution of problems in the Internet-specific aspects of the transaction. It seems unlikely that Internet commerce can flourish without such an intermediary. While it is certainly conceptually possible that such services could be provided by existing financial institutions, it must be

---

[2] Over time, established brand-name identities may help reassure customers in such situations, but this is itself problematic. Brand identity in cyberspace may be too-easily damaged by technical circumstances beyond the control of the identified corporation. Moreover, the establishment of brand identities will be in opposition the egalitarian tendencies of the Internet, which will tend to promote small entrepreneurs or "micro-merchants". Finally, anyone with an established brand identity needs to worry a good deal, on the Internet, about imposters speaking in their name.

remembered that the resolution of these problems can be quite complicated technically. Debugging obscure problems with incompatible implementations of Internet protocols is not a core competence of most financial institutions.

By analogy, people rarely object to the role played, in modern commerce, by parcel delivery services and telephone companies. If the Internet were somehow centrally administered, then the Internet-specific aspects of financial transactions would be handled by that central administration in a manner that paralleled the worlds of telephone and parcel services. However, the anarchic nature of the Internet leaves it without any central authority to resolve technical issues that pit buyers against sellers, and these are of paramount importance to the conduct of commerce. Therefore some kind of Internet service bureau seems essential for investigation and resolution of these problems.

To make all of this more concrete, a few examples are given briefly below. The First Virtual team has encountered dozens, perhaps hundreds of these situations, many of them caused by "sophisticated" multinational corporations, and sees no likelihood that they will stop arising in the foreseeable future. Each new Internet software package or site seems to introduce new bugs arising from incompatible protocol implementation and the like, and ALL of these have an inevitable effect on the conduct of commerce. A few selected examples:

FTP bugs: Some browser software puts an arbitrarily low maximum size on ftp file transfers. The net result is that the buyer gets a truncated file, which is often useless (e.g. for software). However, the seller believes that the buyer has successfully downloaded the software, and sends a bill through First Virtual. (Sometimes, the seller should have been able to tell that the download was aborted, but sometimes this is impossible.) This problem was first introduced when a Fortune 500 computer company began selling products using First Virtual, which demonstrates that technical sophistication is no protection.

Connectivity glitches: Sometimes a partial Internet outage occurs after a buyer has paid for access to a site, but before he or she has been able to reap the benefit of it. From the buyer perspective, this looks like an attempt to "take the money and run", but it really often turns out to be the part of what is essentially a SEVENTH party -- neither the buyer, the seller, First Virtual, nor the ISP (Internet Service Provider) of any of those three, but instead an ISP who is used "invisibly" to connect the other ISP's to each other.

Catastrophic failures on the seller end: If a site sells subscriptions, and then has a catastrophic hardware failure, they are often unable even to tell their customers about the problem. Naturally, the paying customers feel the need to complain to someone and perhaps seek a refund.

Protocol violations: There are many well-known software vendors that provide broken implementations of core Internet protocols. Merchants that seek to make use of some of the higher-end features of the Internet are quite likely to encounter customers whose software doesn't work right. From the customer's perspective,

it's difficult not to blame a merchant who promised a daily picture delivery by email, if the customer sees only a daily message that appears to be garbage (because of a broken MIME implementation, for example). Such bugs are far from rare -- they are found in widely-used software from such well-known providers as Lotus and Netcom, for example.

Unanticipated email limitations: Any services that sell information by email, or particularly that provide email-activated robots, are likely to encounter problems with software that imposes arbitrary limitations. For example, the Prodigy system truncates email Subject headers to an extremely short length, which messes up many robots that key off the subject headers, leaving the Prodigy customers feeling cheated when they don't get a proper response.

Unidirectional communication: Many merchants attract customers to their Web pages, where they ask the customer for an email address. They then use the email address either to deliver information products or vital information about how to use or upgrade products that are delivered physically. Unfortunately, nearly half of all Internet users make a mistake when asked to type in their email address, and thus provide an address that does not work. When they then don't hear from the merchant, they feel cheated. But, of course, it is impossible for the merchant to communicate anything to the buyer without the buyer's real email address, and sometimes this even means it is impossible to deliver the merchant's product. (This problem was vividly demonstrated when the Academy Awards web site tried to run a contest, but couldn't communicate with a third of the entrants because they had mis-typed their email addresses.)

Software configuration bugs: The widely-used Netscape browser, for example, can be used to send mail. When users send mail with Netscape, the mail contains a reply address that is entered by the user as part of a configuration screen. Unfortunately, Netscape makes no attempt to make sure that the address is valid (or even syntactically legal!), so many Netscape users send mail, never receive a reply, and are left with a lack of confidence in either the mail recipient, the Internet mail system, or both, never suspecting that their own browser and/or its configuration are really to blame.

The above examples are used for illustration only; the actual number of such problems appears to be, for all practical purposes, without limit. Each major new service that comes on line seems to exhibit at least one of these bugs, at least for a while. (The recently-released Microsoft Network exhibits almost all of them, and more!) As long as the Internet is full of such glitches, there must inevitably be some kind of Internet-based intermediary for commercial transactions conducted via the net. In order to resolve these situations, the intermediary must have a very deep understanding of the way the Internet protocols actually work. In the last year, First Virtual's team has come to supplement that deep understanding with hundreds of detailed examples, most of which are reflected in patches to the system that work around other peoples' bugs.

In the long term, it is important for the Internet community to achieve a much greater degree of interworking between applications at the highest levels. Internet commerce will increase the demands of Internet users for service providers to provide software that works with everyone else's software, instead of application software that includes so-called "features" that do not interwork with other software. First Virtual believes that market demand for interworking applications will in due course persuade all Internet software vendors to more closely adhere to the open IETF standards. For now, however, there are substantial problems of interoperability and confusion caused by vendors trying to unilaterally define or extend the standards for Internet applications. First Virtual's interim strategy is to simultaneously work around, or "patch", the current problems, and to exert pressure for conformance on non-conforming service providers and application vendors.

C. Security and Administrative Issues

The importance of Internet site security is widely discussed and well-understood. It is of particular importance, of course, in the operation of a commerce server, as such a server is an obvious prime target for would-be criminals. First Virtual began with the assumption that our success would invite ever more frequent and more serious criminal attacks.

There is no reason to doubt that assumption. Our monitoring software reveals regular break-in attempts from various sites, although none, to our knowledge, have succeeded. While we choose not to document, in a public paper such as this one, the techniques and tools that we use to thwart such break-ins, we wish to stress that the risks involved are very real and very serious. Anyone contemplating the implementation of an Internet commerce server should not only acquire significant in-house expertise on Internet security, but should also regularly hire outside teams to test that security and report any flaws found. The same teams should not be used repeatedly, as they will exhaust their bags of tricks before long.

Unfortunately, the more secure you make your server, the more difficult it is to administer it, especially remotely. Even for a commerce system based on non-cryptographic mechanisms, such as First Virtual's, cryptographic tools are essential for secure remote access to the server. (In fact, First Virtual commissioned the development of PGP-encrypted telnet for just this purpose.) Special attention should be paid to the issue of the lifetime of cryptographic keys, as discussed in the section on cryptography later in this paper.

While this section is necessarily short on details, there is a very clear lesson that should be understood by anyone with sensitive information on an Internet connected machine: there are many criminals out there, and they *will* try to break in, either for financial gain or for sport. You must inconvenience yourself to a considerable degree, and at considerable expense, if you want to thwart them.

D. Customer Service Issues

Beyond the previously-discussed need for an Internet intermediary, running a commerce system on the Internet entails a host of customer service issues that may not be obvious at first glance, especially to those already extremely comfortable with life on-line.

It has been pointed out that, because the Internet population doubles every 11 to 13 months or so, at any given moment more than half the user community has been on the net for less than a year. In other words, "newbies" are the rule, not the exception. The reality is that an ever-increasing proportion of the Internet's population has only the barest, most rudimentary understanding of how anything on the Internet -- or on their computer -- actually works.

Compounding this is the ever-increasing number of Internet users whose command of the English language is quite limited. Although English is often described as the de facto language of computing and the Internet, this is neither a completely accurate description nor one that sits well with members of other linguistic communities. Internet commerce systems are inevitably international, and when a customer in Japan buys from a vendor in Japan, it is unreasonable to assume that both will be fluent in English if they need to discuss a problem with the transaction.

The combination of poor Internet understanding, questionable English skills, and real money on the line often creates a confrontational situation. While some problems occur due to actual bugs in the commerce system, the vast majority are some form of "pilot error" or are due to Internet problems outside the domain of the commerce system. It therefore seems likely that the customer service load is for the most part not a consequence of our server design, and must be factored in to virtually any plan to provide Internet commerce services. (Indeed, cryptographically based schemes, which entail the provision of public key technology to naive users, are likely to carry an even heavier customer support load.)

Although we have tried very hard, First Virtual has not always been commended for the timeliness of its customer service. The application domain is very new, the questions very numerous, and the user base doubled every six weeks for most of the first year. On several occasions, the help department has become seriously backlogged. We would recommend that anyone contemplating a similar service should plan on excess capacity in their customer support department. On the positive side, however, is our observation that a sizable majority of all customer support interactions are with new customers in their first few interactions with the system or with the Internet. Once users are familiar with the system, they ask relatively few questions, and the questions asked by new users generally come down to a few common issues which are easily answered, often resolvable with further automation, and which should become less common as the system's documentation continues to improve.

E. Cryptography: Myths and Realities

One of the most misunderstood aspects of Internet commerce is the role of cryptography. Some parties have claimed that safe commerce is impossible without cryptography. Others have (incorrectly) interpreted First Virtual's non-cryptographic system as evidence that our company is philosophically opposed to the use of cryptography. Discussions of the issue have often been unclear and ill-informed. Not surprisingly, we have given these issues a great deal of thought in recent months, and have reached some tentative conclusions.

By "cryptographic technology" we refer here primarily to technology that incorporates the algorithms known as "public key cryptography". In public key technology, each user has two keys, a public key and a secret key. The public key may be handed out far and wide, while the secret key is guarded as closely as possible by the user with whom it is associated. To the extent that one possesses a public key, and believes that this public key is associated with a given user, and further believes that this user's secret key has not been obtained by any third party, one can use public key cryptography for either authentication, encryption, or both.

Of course, the last sentence points to some critical perils of cryptography. Sometimes, a secret key will be stolen without the knowledge of the user with whom it is associated. Other times, a public key that is supposed to belong to a given user may be illicitly replaced by a public key belonging to a third party. Either of these events will completely undermine the utility of the cryptographic algorithms. Thus, a safe application of cryptographic technology will pay close attention to how public keys are associated with user identities, how stolen keys are detected and revoked, and how long a stolen key is useful to a criminal. It is beyond the scope of this paper to discuss the infrastructure and customer support requirements involved in providing and authenticating cryptographic keys for each of the world's credit cards, which number in the hundreds of millions, but we're not sure it can be done at all.

A major factor that can limit these risks is the notion of key lifetimes, in which a public/secret key pair is explicitly declared in advance to be useful only until a certain date. Ultimately, most possible attacks against the overall security of a cryptographic system come down to the issue of undermining the trustworthiness of cryptographic keys. The longer-lived the keys are, the more likely it is that some such attack will undermine their value. Therefore, it is vital that systems which rely heavily on cryptography should bound the lifetime of their keys, so that the consequences of any key-based attack are limited to the lifespan of those keys. Just as your Internet Service provider urges you to change your password regularly, so too your cryptographic software provider should encourage you to change your keys regularly. This is an area with crucial security consequences which are often totally neglected by proponents of cryptographic solutions. People routinely ask, when comparing cryptographic solutions, "how many bits long are the keys?" -- a question which refers to the difficulty of a direct computational attack to break the cryptography. A similarly simple question that can be asked about all cryptographic schemes is, "how long-lived are the keys?" For example, a 1024-bit key with a 5-year key lifetime is probably considerably more vulnerable to criminal attack than a 512-bit key with a 1-month key lifetime.

In assessing the importance of the various risks, it is vitally important to distinguish between the two main applications of cryptographic technology: authentication and encryption. These are often confused or conflated, because they both utilize the same underlying cryptographic algorithms, but they are very different and must be discussed separately for a clear understanding.

Cryptographic authentication -- also known as digital signatures -- is the use of cryptographic technology to provide a relatively strong amount of evidence regarding the origin of digital information. If you are confident that you have my public key, and that my secret key has not been stolen, I can use cryptographic authentication technology to give you a relatively strong level of assurance that the information I send you really comes from me.

Encryption, on the other hand, is the use of cryptographic technology to encode information in such a way that it can only be read by certain parties, not by others. It depends primarily on the sender's belief that he has the receiver's public key, and that the receiver's secret key has not been stolen.

It is important to distinguish between authentication and encryption because these two uses of cryptographic technology have radically different implications in commerce systems, at both the legal and technical levels. Legally, nearly all of the problematic restrictions apply to encryption, not authentication, because governments are concerned about being able to detect spying and other criminal activity. This may or may not be an issue in commerce.

Technically, however, the differences between authentication and encryption are fundamental, and are crucial to commerce in the event that the cryptographic technology is ever compromised or "broken". It is vital to understand the consequences of "broken" cryptography, because no program or algorithm is ever perfect or unbreakable. A realistic analysis of any cryptographic commerce mechanism must therefore include an analysis of the consequences if a malicious party manages to break the cryptography. By "breaking" the cryptography, we refer to either defeating the basic cryptographic algorithms, or to stealing the secret keys involved, or to finding a suitable bug in the software that implements the cryptographic algorithms which can yield the same effect as breaking the overall algorithm, but only for a particular implementation.

In the case of authentication, a criminal who has broken the cryptography can impersonate one or more users. On the Internet, it is fairly easy for the impersonator to make himself completely untraceable. This is obviously a problem, but it is a bounded problem, in that the possible damage caused by the impersonator can be limited. In particular, if someone explicitly claims, on the net, to be Bill Gates, then this allows him only to take those actions that are permitted to Bill Gates. Merchants can limit risk by only allowing Bill Gates to have merchandise delivered to his own home, or can use other methods (such as email or telephone confirmation, for example) to confirm the cryptographically-asserted identity, particularly in the event that the compromise of such authentication has become commonplace.

Encryption, on the other hand, is often more of an all-or-nothing technology. The key to assessing the value of compromising encryption technology is an assessment of the value of the information being encrypted. In the case where a criminal has broken an encryption mechanism, that criminal can read all the encrypted information. Again, the criminal can take steps to be essentially untraceable when he is reading the encrypted information via the Internet. It should be understood that, unlike the authentication case, *there is no fundamental limitation to the cost of having your encryption technology compromised*. The cost of such a criminal act is precisely proportional to the value of the encrypted information. The more valuable your information -- and thus the more likely you are to want to encrypt it -- the less acceptable is the risk of having it stolen by an anonymous malicious party on the Internet. To put it more simply: if information is so valuable that you need to encrypt it, it's almost certainly too dangerous for you to accept the risk of putting it on the Internet in encrypted form, and having that encryption broken. (Note that such considerations apply exclusively to the use of encryption to protect economic value, as opposed to the use of encryption for privacy, which is a very different matter.)

In the case of credit card numbers, the information most commonly proposed for encryption on the Internet, the logic is simple. Imagine a world in which millions of credit card transactions travel over the Internet, encrypted, every day. If a malicious party finds a flaw that allows him to decrypt that traffic, he has now untraceably obtained a stream of credit card numbers that is, for all intents and purposes, infinite. While the credit card system has evolved to tolerate a certain rate of fraud, it is unlikely to prosper in a scenario where a single criminal can steal so many card numbers. (This is because credit card fraud today is often traced by a pattern of use and abuse, but a smart criminal who stole millions of cards would only use each once, and would thus be far harder to track down.) If the criminal was truly malicious, and was motivated more by vandalism than raw greed, he could quite conceivably defraud a significant percentage of the world's credit cards in a single day, essentially destroying the integrity of the whole credit card system.

In assessing these risks, it should be understood that the credit card and ATM industries are based on closed networks. The Internet is the most open networking environment imaginable, was not designed with the kinds of safeguards that are taken for granted on closed networks, and allows anyone in the world to gain essentially anonymous access. This is an environment in which the bank card industry has virtually no experience or expertise. Cryptographic solutions are actually much more useful in closed networks than open ones, because they constitute only a part of the overall security (notably, privacy protection against competitive financial institutions) rather than the sole defense against criminals.

The bottom line on cryptography is that it isn't magic, it isn't perfect, and it will always be breakable. A cryptographic system will only be as strong as its weakest link, and one rarely knows in advance what the weakest link will turn out to be. This means, for example, that it doesn't matter how strong your encryption algorithm might be if it is possible to steal the data before it ever gets encrypted, for example via a key management virus that attaches itself to the user's computer and monitors the user's raw

keystrokes. (This is only one of several possible forms of viral attack on cryptographic systems; other possibilities include viruses that replace one key with another, and viruses that intercept dat after it is decrypted.) Similarly, the best encryption in the world is useless if the data can be stolen after it is decrypted, for example by a conventional "break-in" attack on the machine of an Internet-connected merchant, processor, or bank.

An obvious but often-ignored corollary of this bottom line is that, in an Internet commerce system, cryptography should not be permitted to become a critical-path component with a catastrophic cost of failure. This strongly implies, for example, that a partial reliance on cryptographic authentication is far more defensible than a total reliance on cryptographic encryption. While there is undoubtedly a role for encryption technology, it is far better to keep the most valuable information -- including credit card numbers and other sensitive financial instruments -- entirely off the Internet.

A safer approach that still utilizes encryption, for those who really want it, is to design a financial instrument that creates less damaging consequences in the event of theft, such as the Virtual PIN, and then encrypt *that*.

In short, cryptography is a useful tool but not a panacea. Cryptographic encryption, in particular, is often misused by applying it to the encryption of data that would be better kept off-net in the first place. When evaluating a real or proposed use of cryptographic technology, therefore, two key questions should be asked:

-- What are the potential consequences if the cryptographic software is compromised? In the case of many proposed Internet commerce systems, the consequences may be catastrophic.

-- What steps have been taken to bound the risk if the software is compromised? Proponents of cryptography often equate "risk management" with "key length", on the (mathematically sound) theory that the larger the encryption keys, the harder it is to break the cryptography. This is only true, however, for "head-on" attacks that try to break the cryptographic algorithms themselves. Key lengths are totally irrelevant to attacks based on stolen keys or software bugs. Limited key lifetimes are a much more effective tool in bounding the consequences of "back-door" attacks, but many proposed Internet commerce systems have key lifetimes that are indefinite or unspecified -- that is, the keys last forever. The bottom line is that every system has risks, and it is crucial to understand the manner and consequences of a catastrophic failure.

Overall, First Virtual’s experience with running a completely non-cryptographic payment system has been highly positive, with fraud rates so low as to elicit the excited attention of banking partners. This does not imply that the First Virtual system will forever remain non-cryptographic; indeed, the limited use of cryptographic authentication is being implemented for First Virtual’s second system as of this writing. (And in answer to the questions that should always be asked about such systems: First Virtual will be using 1024-bit keys with 1-month key lifetimes.) However, First Virtual’s experience strongly suggests that cryptography is at most a single tool in the pursuit of security, and is neither an absolute requirement nor the panacea that its proponents often suggest.

To be clear, First Virtual will remain very different from the cryptographic systems being discussed. The non-encrypted Virtual PIN will remain the basis of the system. However, there will be optional facilities for encrypted communication between First Virtual and its more sophisticated merchants, for purposes of hard goods fulfillment and delivery. The problem of deploying cryptographic keys to buyers, and managing all those keys, will therefore not be faced by First Virtual.

IV. Where Are We Going?

After one year of operation, First Virtual's biggest problem is clearly growth management. With a user base and transaction volume doubling every six weeks, we face significant operational challenges. As of this writing, the growth had helped cause one significant operational outage (in August), and that outage attracted wide publicity and concern. Naturally, First Virtual has been devoting a great deal of effort to trying to avoid any further such outages. (Actually, we knew all along precisely what measures could have prevented or minimized this kind of outage, but had held back due to constraints in both manpower and capital. Such tradeoffs are faced by any new business, and sometimes the choices made don't look so smart in retrospect.)

Beyond the struggle to simply provide good service in the face of such growth, however, the First Virtual system is being expanded in multiple directions. At approximately the time this paper is published, the system is expected to be upgraded to better permit the sale of physical goods and services, as opposed to the information products for which the system was originally designed. These enhancements will include the use of cryptographic authentication of certain critical messages sent from First Virtual to our merchants. Future enhancements will include internationalization (for languages and currencies), additional mechanisms for buyers to pay into the system and for sellers to receive payment, and better support for extremely small transactions, sometimes known as "micropayments". Another priority is to open the system to participation by multiple processors and acquirers in the banking world.

A brief mention should be made about why the initial First Virtual system was limited to information products, as opposed to physical goods. The answer is twofold. First, we were enamored with the unique aspects of information commerce, and the consideration of this uniqueness was what led to the initial design of our system. Second, although First Virtual is a pioneering company, it is also a conservative one, with conservative founders and backers. The risk involved in any loss is far higher for those selling physical goods, and it was appealing to "shake down" the system before encouraging anyone to depend on it for such applications. The lessons of that shakedown period, as presented in this paper, have guided the development of additional mechanisms that we believe will make the system completely suitable for commerce in physical goods.

In the larger world of Internet commerce, we expect that there will be a gradual sorting out of the issues, as the nature of Internet commerce becomes clearer. We expect to see, at a minimum, a growing realization that there must be some kind of Internet-based intermediary to help facilitate the technical aspects of Internet commerce. As far as cryptography is concerned, there will probably be a continuing series of "scandals" as it

becomes clear that no encryption software is unbreakable, and that Internet commerce cannot depend upon the existence of unbreakable encryption. One fear is that this may cause a backlash against cryptography, in which the baby is thrown out with the bathwater, and the many practical benefits of cryptographic technology would fall into disrepute. First Virtual will do what it can to make sure that this does not happen.

V. Conclusions

When First Virtual's system went live on October 15, 1994, there was still widespread skepticism that Internet commerce would ever really take off. A year later, such skepticism has largely vanished, in favor of wild speculation and press release fever about the mechanisms of such commerce. Meanwhile, a few pioneers have actually been doing business in cyberspace, making some money and encountering some unexpected problems and misconceptions.

The biggest unexpected problems center around customer service. The Internet is a complicated place, and it isn't getting any simpler. An Internet-savvy customer service department is an absolute prerequisite for anyone providing commerce services to the net.

The biggest misconception is that the words "security" and "encryption" are synonymous, or even closely related. A more balanced perspective on discussions of Internet commerce can often be obtained by replacing "computer" and "encryption" with "automobile" and "door lock". The mere existence of a door lock does not imply that the ignition keys (or a wallet) should be left inside the car. In general, it is safest to lock your car *and* remove your valuables. Similarly, while encryption can provide a modicum of additional security on the Internet, it is far more important to consider what is being encrypted, and not to encrypt anything that is better kept off the net in the first place.

Internet commerce is real, and it is growing at breakneck speed. Early speculations about it have often proven to be far from the mark. The history of the Internet suggests that those who want to play a role in its evolution should start with simple technologies that really work, and expand them from there as circumstances require. First Virtual's initial payment system is clearly only one step in a larger evolution. There are very exciting times ahead.

Bibliography

The best source of basic information about First Virtual's Internet Payment System is the First Virtual Web site, at http://www.fv.com. Most of the same information is also available via mailserver, starting with info@fv.com.

Technical details about the First Virtual payment protocols have been published as Internet Drafts, and will be published as Informational RFC's. They are available for anonymous file transfer from the machine ftp.fv.com, in the directory pub/docs.

[1] Stein, L.H., E.A. Stefferud, N.S. Borenstein, and M.T. Rose, "The Green Commerce Model", First Virtual Holdings Incorporated, June, 1995. File name: pub/docs/green-model.{txt,ps}

[2] Borenstein, N.S., and M.T. Rose, "The application/green-commerce MIME Content-type", First Virtual Holdings Incorporated, June, 1995. File name: pub/docs/agc-spec.{txt,ps}

[3] Rose, M.T., and N.S. Borenstein, "The Simple MIME eXchange Protocol (SMXP)", First Virtual Holdings Incorporated, June, 1995. File name: pub/docs/smxp-spec.{txt,ps}

Those without prior familiarity with the MIME protocol may find the MIME specification invaluable in understanding some of the above documents:

[4] Borenstein, N., and N. Freed, "MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies", RFC 1521, Bellcore, Innosoft, September, 1993.

Appendix A -- Responses to Some Common Questions About First Virtual

First Virtual's Internet Payment System is very unusual, and in many ways flies in the face of conventional thinking about Internet commerce. After a year of successful and secure operation, First Virtual has heard and responded to a wide variety of criticisms, most of them based on a fundamental misunderstanding of one of three things: how the Internet works, how First Virtual works, or how the other approaches to Internet commerce work (or, more often, are purported to work if and when they are deployed). In what follows, we have attempted to collect and respond to 29 of the most common questions about the First Virtual system:

1. Isn't First Virtual vulnerable to sniffing Virtual PINs (First Virtual Account-ID's)?
2. Isn't First Virtual vulnerable to IP spoofing?
3. Isn't First Virtual inherently insecure because Internet email is insecure?
4. Isn't First Virtual inherently unreliable because Internet email is unreliable?
5. Isn't First Virtual Overly Centered on Email? Isn't the Web is Where It's At?
6. Won't First Virtual be made obsolete by the wide-spread deployment of standardized encryption mechanisms?
7. Isn't First Virtual unsuitable for hard goods?
8. Isn't First Virtual unsuitable for merchants with cash flow concerns?
9. Isn't First Virtual unsuitable for very small-cost items (microtransactions)?

10. Isn’t First Virtual unscalable because of its single-acquirer (single bank) model?
11. Don’t merchants get ripped off when buyers say "no"?
12. Doesn’t the First Virtual System Hurt Merchants by Reducing Impulse Buying?
13. Aren’t First Virtual’s fees unreasonably high?
14. Isn’t it too hard for a merchant to sell with First Virtual?
15. Isn’t it too confusing for customers to keep track of their purchases?
16. Doesn’t First Virtual require all sales to be essentially "shareware", i.e. payment-optional?
17. Isn’t First Virtual vulnerable to merchant-based fraud?
18. Isn’t First Virtual vulnerable to service outages due to its centralized architecture?
19. Isn’t First Virtual vulnerable to hacking by system administrators?
20. Isn’t First Virtual vulnerable to break-ins?
21. Isn’t First Virtual basically a low-tech, boring system?
22. Isn’t anonymous digital cash the real answer for network payments?
23. Why should buyers worry about credit cards over the net?
24. Does First Virtual really believe its system is foolproof?
25. How can a criminal break First Virtual’s system, and does it matter?
26. If First Virtual isn’t foolproof, why is it any better than using encryption?
27. What are the real risks to worry about in Internet Commerce?
28. What is the real key to making Internet Commerce as safe as possible?
29. What are the costs, risks, and benefits involved in using the First Virtual system?

We don’t mean to imply, in answering these questions, that we believe we have all the answers to everything to do with Internet commerce. No one does. However, we have thought long and hard about the issues involved, and have answered some of these questions so frequently that we are offering them here in the hope that they will serve to raise the general level of discussion of these issues. We’re always looking for feedback that will deepen our understanding and help to make our system more secure.

1 Isn’t First Virtual vulnerable to sniffing Virtual PINs (First Virtual Account-ID’s)?

All traffic that flows over the Internet is vulnerable to sniffing, because of the open, uncontrolled nature of the Internet’s basic technology. Whether or not such sniffing creates a vulnerability depends on the kind of information being sniffed. For example, if one says, on the Internet, "it is cloudy outside", the fact that this statement can be sniffed does not generally create any vulnerability for the person saying it. If, however, that person describes some secret corporate strategies, sniffing can create a vulnerability.

There are two approaches to dealing with vulnerabilities related to having the "wrong" information sniffed. One approach is information hiding, in which one tries to represent the information in a manner that makes it unusable by the party doing the sniffing. This can be done either by steganography, in which the information is "hidden" in a large mass of other information, so as to make it undetectable, or by cryptography, in which the information is mathematically scrambled so as to make it extremely difficult to unscramble without a "secret key". Neither of these information-hiding technologies is foolproof, but both make it considerably more difficult to utilize the information once it has been sniffed. Such information hiding technologies are much better than nothing, if one really needs to use the Internet to transmit extremely sensitive information. However, if the incentive is great enough, there are always ways for determined criminals to break through the barriers created by either of these technologies.

The second approach to avoiding sniffing vulnerabilities is to keep any such sensitive information completely off the Internet in the first place. Obviously, this is the only foolproof method to avoid losing information to sniffers. When it is possible to keep truly sensitive information entirely off the net, it is nearly always preferable to do so.

The latter is the First Virtual approach, and it explains why we created a new kind of identifier, the Virtual PIN (tm), for use on the Internet. Unlike credit card numbers, for example, a Virtual PIN is almost useless in itself, because it is only worth money in conjunction with an email-based challenge and response cycle. Credit card numbers are designed to be "one-way" instruments, where the mere conveyance of the number is sufficient to cause a financial transaction to happen. Such instruments are, by their very nature, extremely vulnerable to sniffing. Rather than concentrate on technologies to hide such instruments from "sniffers", First Virtual began by accepting the inevitable reality of sniffing on the Internet, and instead designed a financial instrument that could be sniffed without any serious negative consequences.

In short, Virtual PINs can be sniffed, but it’s not a serious vulnerability. Someone who sniffs a Virtual PIN can initiate a transaction, but he cannot consummate it without completing the email response loop on behalf of the real account-holder. Sniffing might be one step in a comprehensive attack on First Virtual, but by itself it is nearly useless.

2 Isn’t First Virtual vulnerable to IP spoofing?

Each Internet-connected computer has a dedicated address known as an IP number, which indicates where it is attached to the Internet. The Internet protocols have been designed around the notion that each machine has an IP address, and each IP address points to a known machine.

IP spoofing is a technique in which one Internet-connected machine pretends to be another one, essentially "hijacking" its IP address. It is relatively easy for a criminal to connect to the Internet and use IP spoofing to steal traffic that was intended for another, target machine. In this manner, all traffic -- notably including email traffic -- that is intended to go to one machine can instead be caused to go to the criminal’s machine.

Because such an approach can cause one machine to temporarily masquerade as another, and hence intercept First Virtual's email queries, it has been suggested that First Virtual's protocols are vulnerable to IP spoofing. If a criminal can intercept and respond to such queries from First Virtual, he can indeed cause some financial transactions to be consummated without the approval of the account-holder.

While it is true that IP spoofing can fool First Virtual as well as any other Internet-connected machine, the utility of this technique for criminals is seriously limited. First of all, IP spoofing is a very crude technology; it generally intercepts all IP traffic for the target machine, and thus leads the users of the real machine to see a serious loss of or degradation of their Internet service. For this reason, it is generally detected quite quickly, in which case it can often be traced to the perpetrator. Thus, IP spoofing is really only useful to a criminal when it can be installed for a brief moment and then removed. This is not the case with First Virtual transactions, because the asynchronous (and randomized) timing of the email queries makes it impossible for a criminal to know precisely when to set up the spoofing.

Moreover, IP spoofing by itself is useless to a criminal who doesn't know what machine to spoof. In order to use IP spoofing against a First Virtual account-holder, a criminal must first steal a Virtual PIN (e.g. by sniffing). He must then determine the email address associated with the stolen Virtual PIN, which is not an obvious association and is not always easy to do. Next, he must use IP spoofing to intercept traffic intended for the victim's machine, and must establish the spoofing for long enough to intercept the mail, but briefly enough to minize his risk of detection. By doing all of that, he may be able to consummate a single fraudulent transaction. However, he will have to take another risk by renewing the IP spoofing every time he wants to make another fraudulent transaction. This is a recipe for risky, detectable, traceable, occasional fraud, not for any kind of serious systematic fraud.

In short, yes, IP spoofing may be used to intercept First Virtual's messages, but no, it's not a serious vulnerability. IP spoofing for more than a brief interval creates serious risks of capture for criminals, but is necessary for extended periods of time in a serious attack against First Virtual. IP spoofing might be one step in a comprehensive attack on First Virtual, but by itself it is nearly useless.

3 Isn't First Virtual inherently insecure because Internet email is insecure?

No, First Virtual is inherently insecure because life is insecure. Nothing is certain except death and taxes.

Seriously, this question, though frequently asked, is not meaningful unless the term "insecure" is defined. Virtual PINs were designed for transmission over an "insecure" medium such as Internet mail. Insofar as the insecurity of concern relates to the theft of valuable financial information, First Virtual is secure because such information is never transmitted over the net in the first place.

Insofar as the concern relates to forged mail, the situation is more complex. It is certainly true that it is remarkably easy to forge Internet mail. First Virtual is not seriously affected by this problem, however, because simply forging Internet mail is not enough to cause a fraudulent transaction to take place. When First Virtual receives an email response authorizing a transaction, that response has to include a special one-time identifier that First Virtual included in the query message it sent to the buyer. Thus, the response mail not only needs to be forged, it needs to include one-time information that had to be obtained from the query.

In short, the ease with which Internet mail can be forged was well-known to the designers of the First Virtual system, and the system makes no reliance at all on the authenticity of the purported sender's address. Forged email might be one step in a comprehensive attack on First Virtual, but by itself it is nearly useless.

4 Isn't First Virtual inherently unreliable because Internet email is unreliable?

Internet mail is not perfectly reliable. Sometimes mail is lost or misdelivered. Sometimes old mail addresses cease to work. Sometimes mail is so mangled by mail gateways that it is useless on the receiving end.

First Virtual's technical team is made up of world-class email experts who well understand these problems. To begin with, we make sure that our customers have a functional email connection before they are assigned a a Virtual PIN. Once they have it, they need to reply to our email queries, or their account will be suspended. Because of the inherent unreliability of Internet mail, we re-send all our queries a few times if they go unanswered for too long, and we have procedures that cause accounts to cease to be valid if their email regularly bounces or is not responded to.

In short, First Virtual achieves reliable operation through careful management of the exceptional conditions created by the occasional unreliability of Internet mail and other Internet facilities.

5 Isn't First Virtual Overly Centered on Email? Isn't the Web is Where It's At?

It's certainly notable that many proposed payment and encryption mechanisms are entirely tied to a single Internet protocol, most commonly the World Wide Web protocols. While we certainly share the global enthusiasm about the Web -- and while our payment system works wonderfully with the Web -- we're skeptical that it's the only mechanism that will be useful for buying and selling on the Internet. In fact, we have already seen quite substantial volumes of First Virtual transactions that were consummated using email, FTP, or Internet Relay Chat instead of the Web. A payment system that works only on the Web is, in our opinion, unreasonably crippled from the outset. For example, how can a merchant reach out to a user who previously visited his Web site to offer him new "sale" merchandise? Many aspects of commerce are by nature asynchronous, while the essence of the web is that users get real-time, or synchronous, response to their actions. In our opinion, email is the natural technology for such many kinds of commercial and marketing efforts.

Moreover, email is the Internet's lowest common denominator. The fact that First Virtual's protocols may be used for payment by email means that even those with less than full Internet access -- indeed, even those with slow, sporadic dialup connections via UUCP or FIDOnet, for example -- can participate in Internet commerce. Moreover, because our only requirement for buyers is that they have Internet-connected email, and because everyone on the Internet has a pre-existing motivation to keep their email service working, many of our customer connectivity problems can be described as "self-healing". In contrast, the cryptographic approaches conjure up images of credit card issuing banks having to field support calls about how their cryptographic "wallet" software interacts with the customer's screen saver, and a myriad of similar problems.

6 Won't First Virtual be made obsolete by the wide-spread deployment of standardized encryption mechanisms?

We don't think so. For a number of reasons, we're skeptical that it will ever happen at all, and we believe that the time to deployment will be quite long in any event.

As long-time Internet veterans, we've been hearing that the widespread deployment of encryption was just around the corner for about 15 years. It never seems to happen, for a large number of reasons, most of them related to the standards process, but also related to legal issues such as patent licensing and export restrictions. Eventually, of course, there will be Internet encryption standards, but even then it isn't clear that encryption is the right way to do commerce. Even assuming that all the standardization and legal hurdles are overcome, and encryption starts to be deployed on a large scale, there are many reasons to believe that it will never be widely adopted for commerce.

The use of encryption to hide valuable information is an inherently risky approach for a number of reasons:

1. If there turn out to be any flaws in the encryption algorithms or software, encryption can become essentially worthless overnight. A system with such flaws could provide a criminal with millions of valuable credit card numbers in a single day.

2. Even if the algorithms and software are absolutely perfect in a vendor's software -- a rare occurrence, to say the least -- it is relatively easy for a criminal to distribute alternate software that appears to implement the secure encryption, but actually steals valuable information instead. It is far from clear how consumers are supposed to know which software to trust, and when. The cryptographic approach requires that consumers only type valuable information into "trusted" software, but begs the question of how consumers can know which software to trust. First Virtual offers simpler, more comprehensible advice to consumers: NEVER type your credit card number or other truly valuable information into any Internet-connected computer.

3. Even if the algorithms and software are perfect AND the consumer knows how to be sure that he's using the right version, it isn't hard for criminals to do an "end run" around them. One clever way to do this is to deploy a virus that sits on a consumer's computer and watches for the activation of the "secure" encryption software. By watching the

user’s key strokes as he enters a credit card number, such a virus can steal valuable information before it is encrypted, and thus never has to worry about "unbreakable" encryption.

4. Encryption schemes are fundamentally safe only as long as all parties concerned take the proper precautions to protect their secret encryption keys. This is, fundamentally, a tricky business. No matter how much effort is put into making the software easy to use, it will always be the case that using cryptography safely is considerably more complicated than setting the clock on a VCR. Given that the average consumer has never mastered the latter task, it is reasonable to be skeptical about the ability of consumers to safely handle and protect their cryptographic keys.

Some people have expressed concern that the technical objections here will be irrelevant if large banks and merchants choose to use encrypted credit cards on the Internet despite these objections. Certainly such decisions could muddy the water for some time. However, we believe that the fundamental flaws in the cryptographic approach will inevitably become clear over time. Every time a vendor of cryptographic software has to publicly announce the fix to another catastrophic bug, and then declares their latest software version to be "bug-free", more people will start to wonder about the wisdom of making all of Internet commerce depend on the unbreakability (and unforgeability) of a few algorithms and programs.

7 Isn’t First Virtual unsuitable for hard goods?

Yes, this is true of the first version of our payment system. But we have no intention of keeping it that way.

First Virtual released the first version of its payment system in fully-operational form on October 15, 1994, with no advance hype or publicity. That system was, and is, generally unsuitable for selling physical goods. It was designed for the safe sale of information, and the fact that a few merchants have chosen to use it to sell physical goods does not imply that it was ever intended for that purpose.

Although First Virtual is a pioneering company, it is also a conservative one. Physical goods sales inherently carry higher risk than information goods sales, and it was appealing to "shake down" the system before encouraging anyone to depend on it for hard goods application s. The lessons of that shakedown period have guided the development of upcoming extensions to the system.

First Virtual recognizes that physical goods are a major part of future Internet commerce, and we have no intention of being a "niche" player. We will be releasing enhancements to our system that facilitate the sale of physical goods in the near future, but we prefer not to make any public announcements or promises until the enhancements are operational.

8 Isn’t First Virtual unsuitable for merchants with cash flow concerns?

Yes. But we have no intention of staying that way.

One of First Virtual’s goals, in designing an Internet Payment System, was to have a very low cost of entry, so that anyone in the world could be a buyer or a seller. As such, we were unwilling to restrict sellers by the traditional screening process that applies to credit card merchant accounts, for example. As a matter of principle, and in deference to the egalitarian nature of the Internet, we wanted to let anyone be a merchant. In this case, our principled goal of universal access was supplemented by our practical desire to serve as the virtual acquirer for millions of new "micro-merchants".

However, there are very good reasons why acquiring banks are selective about merchants. An unscrupulous merchant can perpetrate all manner of fraudulent transactions. Therefore, to protect ourselves and our associated banks, our initial system includes a 91-day hold on all funds to be disbursed to merchants. This avoided risks associated with regulation Z in the United States, and eliminated the need to build a screening procedure for all First Virtual merchants.

This holding period is obviously problematic in the case of established merchants with the cash flow concerns that come from the cost of inventory and production. However, it is perfectly acceptable to a wide variety of merchants selling information products, in which the incremental cost of each sale is virtually zero. These were the kind of merchants targeted in the initial deployment of First Virtual’s Internet Payment System.

Not surprisingly, the 91-day hold has not proven to be the most popular feature of the First Virtual system, and it seems natural to waive it for merchants who do not pose a substantial risk. Application, underwriting, and technical procedures for such a waiver have been under development for some time, and will be deployed as part of the upcoming extension to the First Virtual system to support the sale of hard goods.

9 Isn’t First Virtual unsuitable for very small-cost items (microtransactions)?

No, you can process arbitrarily small transactions with First Virtual, but it’s a bit more work on the merchant end, as you have to accumulate several smaller transactions into one larger one.

The term "microtransactions" turns out to be one that is defined largely in the negative, like "Artificial Intelligence". Artificial Intelligence is basically the study of how to get computers to do anything they can’t currently do. (In fact, the FORTRAN programming language was an outgrowth of early Artificial Intelligence research.) Similarly the term "microtransactions" generally refers to any transactions that are currently too small to handle. Some people in the credit card industry refer to anything under a dollar as a microtransaction, for example.

Thus, by definition, First Virtual doesn't do microtransactions, because we tend to think of microtransactions as anything too small for our current system. Right now, in our system there's a minimum transaction fee of 30 cents, which means that 31 cents is the smallest transaction you can put through the system at all, and 50 cents is the practical minimum price we're actually seeing used. So we tend to think of microtransactions as being anything under 50 cents.

We're very interested in moving that minimum down, of course, but we believe there are fundamental reasons why a centralized payment system will never be ideally situated to do the very smallest of transactions. That's because the basic cost of operation increases with the use of a central server that has extremely high requirements for availability and reliability.

To put it another way, it's fundamentally cheaper for our merchants to add up penny transactions until they reach 50 cents than it is for us to add up those same transactions. That's because the merchant is freer to trade off the cost of redundancy versus the risk of lost data. We can't tolerate any lost data, ever, but a merchant who is accumulating penny transactions might be willing to tolerate the loss of a few hours worth of pennies in the event of a catastrophic disk failure. Thus, ultimately, the very lowest-end transactions will always be more efficiently aggregated in a distributed fashion by our merchants than at our central server.

For this reason, we strongly encourage merchant-level accumulation for any merchant who is in a business where very small transactions make sense. In merchant-level accumulation, the merchant validates the buyer's Virtual PIN when the sale is made, but accumulates these transactions for each Virtual PIN until a threshold is reached in either total amount or elapsed time. The merchant then submits a single First Virtual transaction for the accumulated amount. In such cases, the merchant is also collapsing the risk of all these smaller transactions into a single larger one, for which reason most merchants will also want to put an upper bound on the accumulation function.

While we will continue to try to reduce our costs and lower the minimum transaction amount that can be passed through the system directly, we strongly believe that ANY centralized payment system will have more per-transaction overhead than there is in the case of merchant-level accumulation. For this reason, we encourage all sellers of very-low-cost items to consider this approach. At some point, First Virtual hopes to make software available to help streamline this approach for merchants, but it is not currently a high priority.

10 Isn't First Virtual unscalable because of its single-acquirer (single bank) model?

No. There is nothing inherent in our model that limits it to a single acquiring bank. However, we have a very good relationship with our acquirer, First USA, and we have seen no urgent need so far to involve another one. On the other hand, we recognize that there are lots of banks out there, and that over time many of them will want to process credit card transactions over the Internet. We're not terribly worried yet about outgrowing our bank, but we are confident that we can easily extend the system to

multiple acquiring banks if ever we need to do so.

11 Don’t merchants get ripped off when buyers say "no"?

Not any more than they want to be. What First Virtual has invented is something that is really quite new in the world of electronic payments, but which reflects common practice in the real world: a multi-stage transaction.

When you contract for certain kinds of services -- electricians, plumbers, and so on -- you often don’t pay for the services until after they’re rendered. These service providers occasionally don’t get paid, but in general their risk is low if they know where you live, you haven’t refused to pay them before, and so on. Similarly, the process of buying clothes or books often involves some kind of trial usage in the store, but the store has implicitly screened you by allowing you into the dressing room (and indeed their security staff will quickly eject known criminals). In larger transactions, such as purchasing an automobile, the parties will usually quite explicitly "check each other out" before consummating the transaction -- the buyer will try out the merchandise, and the seller will check out the buyer’s credit history. Economic transactions don’t always happen instantaneously, though many payment mechanisms are built on the assumption that they do.

First Virtual’s mechanism explicitly recognizes that economic transactions happen in several stages, and was particularly designed for the special characteristics of transactions involving the sale of information products. In this model, stage one is the buyer authentication, in which the buyer demonstrates reasonably conclusively to the seller that he is a possible buyer -- in our case, by telling the seller his Virtual PIN, which can be validated by the merchant in real time. In stage two, the seller sends the information to the buyer and the bill to First Virtual. In stage three, the buyer can say yes or no after examining the information product, although a buyer who says "no" too frequently will have his account invalidated.

Some merchants interpret this possibility of a "no" response as a possibility of being "ripped off". However, it is important to note that with pure information products, a buyer can’t really evaluate the product until he has the information in his possession. Forcing the buyer to pay for the information before he can examine it is a recipe for increasing the number of dissatisfied customers. A comparable strategy with clothing, for example, would be to keep all of the clothing under glass, and to forbid customers from trying anything on before they have already paid for it. A store that took this strategy would certainly have to hire more people to staff its returned goods department!

We prefer to think in terms of what we call "intelligent lost sales". Most of those who complain about the possibility that the buyer might say "no" have relatively little experience with the headaches and expenses involved in refunds and chargebacks. If it were possible to tell in advance that a given transaction carried an unusually high risk of refund or chargeback, it would be in the merchant’s interest to forego that transaction. Unfortunately, this is generally not possible, but First Virtual provides a new wrinkle for information sales.

In general, physical goods are sold in stores in a manner that permits at least a certain minimal amount of "try before you buy". To the extent that this is not possible (e.g. with mail order), the most successful merchants have been those with extremely liberal return policies. But handling returned goods is a nuisance, and a totally unnecessary one for information products, where the cost of inventory is zero and the seller never really wants the buyer to "return the bits". Allowing the buyer to try the merchandise after authenticating himself, but before paying, is simply a more efficient business practice, in the information industry, than handling refunds to dissatisfied customers. The customer who tries on a sweater and then decides not to buy it is an "intelligent lost sale", far preferable to the merchant to a customer who buys the sweater and then returns it after trying it on for the first time at home.

In general, a merchant is better off forgoing a sale if it has too high a probability of leading to the need for a refund or chargeback, as the cost of processing the refund can easily exceed the profit margin on the product. With physical goods, it is very important to get the goods back, and thus the cost of processing a refund is more often necessary. With pure information products, however, the economics often point in the other direction.

Having said all of that, however, there's nothing in the First Virtual system that prevents merchants from structuring their business so that they don't deliver the purchased product until AFTER the buyer says "yes". Indeed, this is the logical way to use an First Virtual-like system for selling physical goods. However, our observation and experience so far has confirmed this is a poor strategy for selling pure information products. People are much happier with merchants who let them sample the information before answering yes or no to the transfer-query from First Virtual, and abuses are quite rare for most merchants. Indeed, the frequency of customers saying "no" seems to correlate quite closely with the quality of the merchandise.

12 Doesn't the First Virtual System Hurt Merchants by Reducing Impulse Buying?

Insofar as the need to say yes to an email message reduces impulse buying, it is probably true that First Virtual's system slightly reduces the likelihood of impulse buying. In practice, so far it does not appear to be a very significant reduction -- it seems quite clear that First Virtual account-holders do use their Virtual PINs to browse the Web and buy things impulsively.

In the extreme, impulse buying is a mixed blessing anyway. It seems clear that items purchased on a sudden impulse are more likely to be returned or otherwise rejected by the buyer. The concept of "intelligent lost sales" applies here too -- it's not worthwhile for a merchant to accept an impulse purchase if the customer is likely to request a refund later anyway.

13 Aren't First Virtual's fees unreasonably high?

When a merchant sells something through First Virtual, he pays 29 cents plus 2% of the sale price. This fee includes all bank and credit card fees. While there are a few other minor fees (in particular, $10 to set up a seller's account, and $1 each time an aggregated amount is paid out to the seller by direct bank deposit), this is the only cost of any significance to reasonably large-volume sellers.

We think these fees are remarkably low. In most cases, they're only slightly higher (or no higher at all) than what a merchant would pay if he was accepting credit cards directly. But First Virtual provides a whole set of services that aren't included with normal credit card processing.

For starters, the merchant doesn't need to handle his own connection to a credit card acquiring bank. He doesn't have to deal with any of the exceptional conditions involved in credit card authorization, settlement, and chargebacks. The money just shows up in his bank account. All he does is send a message over the Internet to First Virtual to initiate a transaction, and process the messages that come back from First Virtual that tell him how it's going.

Moreover, First Virtual runs a very elaborate (and costly) customer service department. While the intent of our customer service department is simply to help people when they encounter any snags with First Virtual, we inevitably end up dealing with some problems that are really on the merchants' end, and with all the Internet "newbies" who are too naive to understand how to use our merchants' services. While we try to direct merchant-specific questions to the merchants themselves, we inevitably end up doing a lot of customer education about general Internet mechanisms. This customer service function is also factored in to First Virtual's fees.

The truth is that we've set our fee structure remarkably close to costs, in the expectation that we can only make a profit from the payment system business if the volumes are huge. Everyone would always like to pay lower fees, and we're constantly looking for ways to cut costs, but right now they're as low as we've been able to make them.

14 Isn't it too hard for a merchant to sell with First Virtual?

How hard it is to sell with First Virtual depends a lot on what, exactly, you're trying to do.

The core of First Virtual's business, the starting point, is simply a payment system. A merchant can send billing requests to the system, and the system sends back information about the results. This is the core piece of functionality that was missing until First Virtual came along, and without which it's extremely hard to run a business that depends on payments over the Internet.

Now, this raw service, though useful, isn't a "turnkey" solution. Getting a seller's account with the First Virtual payment system isn't a full business setup by itself, any more than getting a Visa or MasterCard merchant account is enough to set up a full business in the physical world. Each merchant still needs software that will generate the

First Virtual transactions, and software that will collect all the response messages from First Virtual, for accounting, tracking, and so on.

There will probably never by a single piece of software that does these things in a manner that will satisfy ALL merchants. Some merchants will want to sell items one-at-a-time from their Web pages. Others will want to sell subscriptions services by email, with email response robots that can be used to purchase additional items. Others will want to sell monthly subscriptions to a whole range of services. On the back end, some merchants will want to feed all the messages that come from First Virtual into a database, while others may want to automatically print them on paper and feed them into a pre-existing manual accounting system.

In short, First Virtual built its payment system to be a genuine "open system". This means, among other things, that anyone can build something that "plugs into it" in a wide number of ways. As a result, after our first year of operation, there are already a variety of vendors selling First Virtual-enabled server software for various purposes, and many more merchants who have customized their own software to talk directly to First Virtual. These supplement or replace the various free merchant tools that First Virtual gives away on our Web site, which are in fact all that is needed in several common situations. (And, of course, smaller-scale sellers can simply use our InfoHaus, which allows you to sell a fairly wide variety of products without even running your own server, although at an additional cost.) Moreover, many more such tools are in the pipeline, both from First Virtual and others.

There are vendors who try to offer a one-stop-shop for everything you need to get onto the Internet. This is a valuable service, but also a dangerous one. Not every customer has the same needs, and such vendors will sometimes try to convince you to change what you think you need in order to suit the capabilities of their products. First Virtual started at the other end -- with a core bit of functionality, the payment system, that we knew everyone would need. This means that merchants have more options (all those third party merchant tools to choose from), and therefore that getting started requires a bit more thought. We think that’s a good thing. We think that merchants setting up new businesses on the net should take a little while to think through their preferred business model to the point where they know what they want their software to do, and then to look for software that does it. We prefer not to try to tell our merchants how to structure their business; that’s up to them. Our primary job is to make sure they can bill customers and collect payments.

Of course, as we grow, and as we learn the applications in which our customers seem to be having difficulty finding or developing appropriate software, we intend to continue to offer more merchant-enabling software, some for free and some commercially. The Internet is still in the very early days of doing real business over the net, and we’re learning new things every day.

15 Isn’t it too confusing for customers to keep track of their purchases?

As is the case with merchants, every customer likes to do things differently. Some will want to feed information about their purchases into a personal financial program on their PC. Others might want a standalone package for accounting for First Virtual transactions. Others might want to ignore the accounting messages entirely. And, given the paper nature of the underlying credit card world, many prefer to print our messages on paper and keep them with their credit card receipts.

All of these are options in our system, though there isn't yet software that easily feeds First Virtual transactions into most personal financial packages. The open nature of First Virtual's specifications makes this easy to do, however, and it seems inevitable that the vendors of such packages will add such support if and when their customers demand it. So far, most customers seem happy with just printing out their "payin-notification" messages, and treating them like normal credit card receipts.

16 Doesn't First Virtual require all sales to be essentially "shareware", i.e. payment-optional?

No. Actually, there are at least three different strategies for selling things with First Virtual. In the most conservative strategy, and the one most suitable for physical goods, the purchased product is not delivered until after the buyer says "yes". In First Virtual's recommended strategy for information sales, the information is delivered after the buyer identifies himself as an First Virtual account-holder in good standing, but before he says "yes" to the transfer-query. This is not the same as shareware, because he has identified himself and knows that his account will be revoked if he says "no" too frequently. In the third scheme, the seller can add a "payment-optional" flag when requesting payment from First Virtual, yielding a true shareware mechanism.

Understanding the meaning of the "payment-optional" mechanism is the key to understanding the difference between First Virtual's recommended mechanism for information sales and true shareware. In shareware, the seller is literally giving away his information, and politely requesting payment from satisfied customers. With normal use of First Virtual's system, the seller is explicitly selling his information, on terms that indicate that payment should be made unless the buyer is dissatisfied, and in the knowledge that strong limits exist to prevent buyers from abusing the system by too-frequently claiming to be dissatisfied. A shareware vendor can set the "payment-optional" flag to indicate that the buyer can say "no" without harm to his First Virtual account history, but this is an choice that the seller can make, not the buyer.

17 Isn't First Virtual vulnerable to merchant-based fraud?

The essential nature of the credit card system is that merchants are reasonably well-trusted, though they are still usually required to provide certain guarantees to their acquiring banks. An unscrupulous but clever individual can do considerably more damage with a credit card merchant account than he can do with a stolen credit card. This is why, in general, getting a credit card merchant account involves a serious screening process, and why it is often a major hurdle for would-be entrepreneurs.

Given the egalitarian nature of the Internet and the beliefs of First Virtual's founders, including the belief that an international industry of micro-merchants was just waiting to be enabled, we built a system in which absolutely anyone could be a merchant. This, on the face of it, opened us up to some spectacular fraud scenarios. However, we then added the 91-day fund holding period, as discussed above. This holding period protects us against nearly all conceivable merchant-based attacks, as they are generally detected within a billing cycle or two. First Virtual has control of the funds for roughly three billing cycles, and earns the interest on the funds during that period.

As we begin to waive the 91-day fund holding period for qualified merchants, we and our acquiring bank will be as vulnerable to fraud from those merchants as an acquiring bank is to fraud from its normal Visa and MasterCard merchants. For this reason, we intend to make sure that our qualification process is complete and rigorous. Companies who can't get a Visa or MasterCard merchant account aren't likely to qualify for waiver of the First Virtual 91-day hold, either.

However, we do have the potential to do something unique here. Someone who can't qualify for a normal merchant account, and who can't qualify for First Virtual's faster payout, can still start a business using First Virtual, as long as they can live with the 91-day holding period. If their business is successful, over time they will develop a history with us that will eventually allow them to qualify to have the holding period waived. In this regard, we think of our system as a "merchant incubator", and we think it's something quite new and much-needed in the realm of financial instruments for would-be entrepreneurs.

18 Isn't First Virtual vulnerable to service outages due to its centralized architecture?

One of the best things about the Internet is that it's so free-wheeling and democratic. One of the worst things about the Internet is also that it's so free-wheeling and democratic. No single organization owns or operates the Internet, and often the resolution of transient technical problems depends on the cooperation of multiple parties, some of whom (the Internet Service Providers) are in direct competition with each other. The system generally works remarkably well, but this isn't a recipe for 100% availability and reliability.

As long as First Virtual uses the Internet to communicate with its merchants, there is the possibility that, at any moment, connectivity will be wholly or partially lost for any number of reasons, many of which are outside of the control of either First Virtual or the affected merchant and consumer. We were very aware of this when we designed our system, and we designed a system that could function in a wide variety of partial or total network outages. However, not every aspect of that design is in place yet, and using it correctly requires a degree of sophistication that not all of our merchants currently have.

The one thing in First Virtual's system that sometimes needs real-time availability is the account validation process. Some merchants need to be able to ask, "Is this a valid First Virtual account" and get a real-time answer. (Note that not all merchants need even this

feature in real-time, e.g. if they're only delivery their product asynchronously anyway.) For this reason, we designed this function to be replicated and distributed to several secure servers on various parts of the Internet. Thus, even if our primary site is no longer accessible to a merchant (which could happen even though both sites are functioning, if some intermediate routing site is having problems), a merchant can answer this question at one of the secondary sites. For this to work, however, merchants have to explicitly try the multiple addresses in the Domain Name System for "inquiry.card.com". Much of the merchant software that is currently out there looks only at "card.com" for this purpose, and is thus more vulnerable to service outages in the event that card.com becomes temporarily inaccessible. Moreover, for our first year of operation the "distributed" functionality wasn't truly distributed, and it behaved more like a monolithic system, all disappearing at once during certain network problems. The full replication is being deployed as this answer is being written.

Another important aspect of maximizing the overall system reliability is to understand that, aside from the account validation step, all of our other transactions are designed to fall back to email. Thus, if your software wants to send a bill to a customer, but our server is inaccessible, your software can simply queue up the transaction via email. Email is the one major Internet service that was specifically designed to work reliably across every kind of network outage, because of its queue-based, store-and-forward nature. Once you validate an account using our replicated validation service, you can send us the actual bill by email, and it doesn't really matter if all the servers at our primary site are in a million pieces on the floor, with a technician standing next to them scratching his head. As long as the technician gets it back together later on, your mail will be received and your transaction will be processed.

Finally, our server reliability and responsiveness can also be enhanced through a variety of techniques that are well-known and used for a variety of network services. For example, we can decrease the frequency of connectivity losses through the use of redundant Internet connections with different Internet Service Providers. We can increase our server performance through the use of multiprocessing systems and by off-loading certain tasks, such mail queue servicing. A shared file system can permit multiple machines to be used to further increase the throughput on the "centralized" server. We do not envision anything more than transient bottlenecks as the system continues to grow and evolve.

In summary, our system was designed to tolerate the transient outages and service problems that are an inherent and permanent part of the Internet. However, not all parts of that design were deployed immediately, and not all merchant software has properly utilized the features that increase reliability during outages. These are problems of evolution, a natural consequence of First Virtual's evolutionary strategy, and are improving over time.

19 Isn't First Virtual vulnerable to hacking by system administrators?

Absolutely, yes. An unscrupulous system administrator can use his access to email accounts on his system to compromise any First Virtual accounts associated with those users' email addresses.

This is hardly ideal, but there are absolutely no better alternatives. Some people believe that cryptographically-based systems don't have this problem, but those people are dead wrong. A system administrator is, by definition, the person who installs "trusted" software on the end user's system. He can easily install what appears to be a new version of the trusted encryption software, but which actually steals all the vital information before encrypting and transmitting it. The bottom line is that your system administrator is really a very powerful person.

System administrators have always been powerful people in the computer world. They can read your private files, read your mail, forge mail from you, fake evidence to make you look like a criminal, and all sorts of other nasty things. Not everyone realizes this, but system administrators need to be highly trustworthy people, just like bankers, notaries, and people who fill a number of other important roles in our society. With the advent of Internet commerce, the role of the system administrator inevitably becomes even more important, but this isn't specific to First Virtual's system. It applies to every Internet commerce system we have heard of or imagined.

On a very large system with many users, it may become necessary, over time, to establish more rigorous screening and qualification criteria for system administrators. A system administrator on a timesharing system like AOL, for example, already has the power to do a lot of mischief. With Internet commerce, he has new opportunities for fraud and theft, as well. It is to be hoped that established Internet Service Providers will generally have strong screening criteria for system administrators.

20 Isn't First Virtual vulnerable to break-ins?

Anyone who connects a machine to the Internet is potentially vulnerable to break-ins. Anyone processing financial transactions is an obvious target. First Virtual was, to our knowledge, the first company to establish a connection (albeit a very indirect one) between the Internet and the US financial networks, including the credit card processing networks and the Federal Reserve. From the beginning, we have regarded Internet break-ins as the most serious security threat we had to worry about.

As a consequence, we have one of the most sophisticated security systems around, with multiple kinds of layered firewalls, watchdog processes, alarm mechanisms, and a rather unusual internal network architecture, the details of which we prefer not to go into in public for the obvious reasons.

None of this, of course, can guarantee absolute protection against break-ins. To further increase our confidence, we regularly employ outside contractors to probe our security and report any weaknesses they might find. We use rotating outsiders for this purpose because we know that different people know different tricks and look for different weaknesses. We expect to continue this practice indefinitely.

In short, we do everything we can think of to run the most secure site on the Internet, but we still can't claim to be invulnerable to break-ins. In our first year of operation, we believe that there have been no successful break-ins, though there have been plenty of attempts.

It's also worth noting that even if someone did break in to our site, they wouldn't find any bankcard information there. All such vital information is kept off-net, on a secure machine at EDS to which our machine is connected by proprietary, non-Internet mechanisms. A criminal who got into our machines would then need to start trying to figure out how to reverse-engineer those mechanisms, which would be a very lengthy process.

In fact, the sensitive data is so carefully shielded that First Virtual can't tell you the credit card number associated with your Virtual PIN even if you ask for it. Only EDS has that information. Some of our customers are skeptical when our customer service department tells them this fact, but it's absolutely true.

21 Isn't First Virtual basically a low-tech, boring system?

This question typically comes from people who are enamored with complex, innovative, and mathematically elegant cryptographic schemes for commerce. Compared with them, we may indeed appear to be low-tech and boring. On the other hand, our system actually works, and this is in part because we designed it in such a way that there's no single "high-tech" component which will make headlines if it turns out to have a single serious bug.

We want to be in the business of processing financial transactions on the Internet for a long, long time. One should be very conservative in the deployment of new technology when vast sums of money are on the line. Bankers are often maligned as boring, too, but they generally take it as a compliment. Where the safety of your money is concerned, we try to be as conservative and boring as possible. That's why, for example, all the money we collect from the bankcard associations is immediately deposited in an escrow account at First USA bank as a further safeguard. We're as boring as we can possibly manage to be while staying on the leading edge of Internet commerce.

It's a little harder for us to take "low-tech" as a compliment, because we're really a company full of high-tech, gadget-happy people. The founders of First Virtual are almost a "Who's Who" of Internet pioneers. And we think that the sophisticated email management-and-response systems that we've built at First Virtual are in fact state-of-the-art and high-tech. The "low-tech" complaint generally comes from people who wish we were using cryptography.

The fact is, however, that we ARE using cryptography -- we use it very heavily for our sensitive internal communications, which are vital to a distributed corporation that communicates mostly by Internet email. Perhaps surprisingly, this experience has only further convinced us of the foolishness of trying to deploy encryption technology to the average consumer. Even very smart non-technical people have a hard time learning how

to use cryptographic software and safely manage their cryptographic keys. Teaching our internal staff to use cryptography safely was enough to convince us we didn't want to try to teach the masses.

There probably aren't very many companies around where a non-technical CEO has learned how to send and receive encrypted email, but we're one of them. When our CEO meets with CEO's from companies that want to use cryptographic technology for Internet commerce, guess which CEO is ready to exchange and validate encryption keys? Guess which one typically doesn't know how to send and receive encrypted email? From our perspective, this hardly fosters confidence in the ability of these companies to safely deploy encryption technology.

Cryptographic technology does have a role in commerce, and we think we know a few things about what that role is. We will be phasing cryptographic technology into our system as an option in the near future. But we probably won't ever be expecting consumers to use cryptography, because we don't think there's any way to teach them to do it safely. If that makes us low-tech and boring, we gladly accept the label.

22 Isn't anonymous digital cash the real answer for network payments?

Anonymous digital cash is a beautiful technology. Its mathematical elegance, for those who understand it, is enough to take one's breath away. It is impossible to read about anonymous digital cash and the underlying algorithms without feeling admiration for the inventors and excitement about the idea. The previous question not withstanding, we are hardly so low-tech as to be immune to this appeal.

For better or worse, however, there are a lot of practical problems with anonymous digital cash. First of all, the very anonymity that excites people turns out not to be all that desirable in most financial transactions. Certainly nobody wants to buy or sell a house anonymously; they'd prefer to know who they were dealing with. On the other extreme, the transactions where anonymity is most fervently desired are often those that are least socially acceptable, which somewhat mutes the enthusiasm many people feel for the technology.

In between, of course, are the vast set of everyday transactions that make up most of our financial lives. Here, the desire for anonymity basically comes down to a question of privacy. I'm not really ashamed of buying my clothes at the Salvation Army (or the Fat Person's Store, or whatever), but neither do I want the whole world to know about it. Unfortunately, this perfectly valid desire for privacy is squarely in conflict with the desire of governments to tax and regulate commerce. While some people see the disempowerment of governments as a feature, rather than a bug, many people worry about depriving the government of the funding it needs to carry out its essential operations (whatever you may believe those to be). In short, anonymous digital cash would be a very mixed blessing in the real world.

First Virtual's system tries hard to straddle the line between the extremes of total anonymity and a total lack of privacy. A Virtual PIN is the only identifier you need to give a merchant in order to buy with First Virtual, and it doesn't reveal anything at all about your real identity. However, First Virtual can, if it needs to, track down your real identity -- even if you lie to us about it, we do have your credit card number, from which you can be traced (unless, of course, you were using a stolen credit card to begin with). We, in turn, maintain a policy of protecting your privacy to the absolute maximum extent permitted by the law.

What this means in practice is that if you want anonymous (or, to be more technically accurate in First Virtual's case, pseudonymous) transactions, you can have them with First Virtual unless the government comes to us with a court order, in which case we'll have to spill the beans. Assuming a reasonable amount of respect for the integrity of the US government (we're a US company, so that's where the court order would have to come from), this should suffice to protect your privacy if you're really not doing anything illegal.

First Virtual as a company does not have a position on anonymous cash as a political/social issue. We are neither for it nor against it. We are open to the possibility of interoperating with it, for example, by using First Virtual to buy or sell such cash, within the limits imposed by the applicable laws and banking and credit card regulations. However, it was not the technology we chose to base our business on, in large part because of the difficulties and vulnerabilities we see in any system that relies too heavily on cryptography, as discussed elsewhere in this document.

23 Why should buyers worry about credit cards over the net?

In the United States, credit card holders are strongly protected, by law, against significant losses due to credit card fraud. In general, the most a consumer can lose, if his credit card is stolen, is $50. Such regulations do not protect consumers in most other countries, however, a fact which is often overlooked in the discussion of the risks of using credit cards in the international Internet.

This protection leads many US consumers to feel little fear about sending credit card numbers over the net, whether they are encrypted or not. What have they got to lose, after all? The answer is simple: $50 and a fair amount of inconvenience. When your credit card number is stolen, you typically end up spending several hours dealing with the credit company's fraud investigators, going at least a few days without your credit card, and all of this can cost you fifty bucks. None of this is the end of the world, but if it happens to you, you're definitely going to be less enthused about sending your credit card number over the net in the future. Moreover, our focus group studies indicate that most people just don't want this kind of aggravation, and will not engage in commerce that they perceive to carry that risk.

In short, the risks to consumers are, in fact, low enough that it is understandable if consumers sometimes choose to take the risk in the absence of any better alternative. But a system where it is so easy to cause that much inconvenience and worry to buyers is a

poor foundation for a long-term Internet commerce system. Moreover, a system where Americans bear little risk but non-Americans bear greater risk is hardly a good starting point for an international commerce system. And, most important, the real risks in sending credit cards over the net are to be found elsewhere, not on the consumer's end, as will be discussed shortly in the section on "the real risks in Internet Commerce".

24 Does First Virtual really believe its system is foolproof?

Not for a moment. There has never been a foolproof system for financial transactions, and there never will be one. Anyone who tells you otherwise is mistaken. Ancient merchants had to learn to judge the authenticity, purity, and quality of gold and other precious metals. Modern merchants have to be alert to forged cash and to credit card fraud. Internet merchants will inevitably have to take certain precautions as well.

First Virtual has always been up-front about the fact that our system can be broken. As discussed above, the exposure involved in such breaks is strongly limited by the system architecture. We will even, in the next section, tell you exactly how to break it! We strongly believe that anyone advocating a system for commerce on the Internet has an obligation to explain exactly how it can be abused, and -- even more important -- what the consequences of such abuse will be. You can be sure that the world's criminals will figure out how to get such information, but you as a merchant have a right to get it first, and you should demand this information from anyone who claims to process payments over the Internet.

The right question to ask is not, "Is the system foolproof?" The only honest answer to that question will always be "no". The right question to ask is, "What are the costs, risks, and benefits involved in using the system?" An honest answer to that question can be very eye-opening. We will answer that question for First Virtual in the final question in this appendix.

25 How can a criminal break First Virtual's system, and does it matter?

To use the First Virtual system to steal things, the best approach we've thought of is described here. There are many minor variations possible, but we believe that this description is representative of the most effective form of successful attack on the system.

Step 1: Steal a Virtual PIN. There are many ways to do this -- Internet "sniffing" is one obvious way, but that particular technique makes step 2 a bit harder. Breaking in to someone's system and spying on them when they type in their Virtual PIN is probably the best approach. (You can sometimes get a Virtual PIN by "shoulder surfing" -- peeking over someone's physical shoulder -- but this approach can make steps 2 and 3 harder, and you have to break into their system at that stage anyway.)

Step 2: Figure out what email address is associated with the Virtual PIN. This can either be fairly easy or very hard, depending on how you got the Virtual PIN. In general, the Virtual PIN and the email address don't usually travel together, but you can sometimes

associate them by traffic analysis or other sniffing-related techniques. Obviously, if you stole the Virtual PIN after breaking into the victim's system, you already know their email address, which makes step 2 trivial.

Step 3: Break in to the computer on which the owner of the Virtual PIN reads his email, and obtain whatever system privileges are necessary to make changes to his email delivery. Note that this step is nearly always a crime in its own right, is often detected, and involves a serious risk of going to jail.

Step 4: Install a mail filtering mechanism so that First Virtual transactions intended for the target user are automatically answered "yes", but so that other email goes through to the user in the normal fashion. This step is mildly tricky to do perfectly and hide completely from the legitimate authorities, but with enough expertise it can be done.

If a criminal successfully executes these steps, he has obtained essentially unlimited use of the victim's credit card number, but only on the Internet, and almost certainly only until the next billing date, at the latest. In other words, for this amount of effort, the criminal has gained a subset (the Internet portion) of the value of actually stealing a single credit card number.

There are three things that should be noted about this description of how to break First Virtual.

First, it is most important to recognize is that this approach is NOT well-suited to automation. Each system break-in is a little bit different, each mail delivery environment is a little bit different, there's no obvious way to automate step #2, and in general it would be very hard for a criminal to perform this kind of crime en masse, millions of times. This means that the overall systemic risks are reasonably well bounded. This is a vital contrast to the vulnerabilities of encryption-based schemes, where an automated attack can yield millions of credit card numbers for a single criminal.

Second, it should be understood that the nature of the risk involved in a financial instrument helps to define its recommended uses. In the physical world, there are many different financial instruments, and they are typically used for different kinds of transactions. The average purchase of bubble gum is made by cash; the average purchase of clothing is made by credit card; the average purchase of an automobile is made by a collateralized loan; the average purchase of a house is made by a mortgage; the average takeover of a large corporation is even more complex. Certainly, it is technically feasible for a merchant to require the same amount of paperwork for a bubble gum purchase as for a corporate takeover, but this would simply convince the consumer that he didn't really want the gum after all. Similarly, we expect that there will be different kinds of mechanisms for Internet commerce, and that different kinds of transactions will be more suited to each. Our mechanism was designed for items in the under-$50 category, although in practice it is already being used up to the $600 range. If it were possible to use First Virtual to transfer millions of dollars to an anonymous Swiss bank account, we would be considerably more concerned about a single occurrence of the kind of attack described above. In a system used for relatively low-ticket items, the big worry is

systematic automated fraud, not occasional fraud.

Third, we must say that the above description of how to break First Virtual is a simplified description and is included for explanation purposes only. In practice, we believe that we would have a good chance of catching anyone who followed the above procedures, and we reserve the right to prosecute such persons to the full extent of the law. The above description in no way constitutes an invitation to try to break or defraud the First Virtual Internet Payment System. In other words: "We are professionals. Do not try this at home."

26 If First Virtual isn't foolproof, why is it any better than using encryption?

There are a number of reasons. If you start from the (false) assumption that the presence of some amount of risk is the same as any other amount of risk, you're still left with the fact that any Internet-connected customer can use First Virtual without obtaining or installing any special software. Thus, First Virtual is preferable on simple grounds of deployment and availability.

However, the more important answer goes to the nature of the risks. Encryption-based systems may be marginally less vulnerable to the kind of casual, one-time attacks that are First Virtual's weakest point, as described in the previous section. However, encryption-based systems have a catastrophic failure mode. If anyone either breaks their encryption algorithms, finds a bug in their encryption software, or does an "end run" around their encryption as described earlier in this document, that criminal can obtain virtually unlimited quantities of valid credit card numbers. The attack is, in general, fully subject to automation, and if the Internet is commonly used for the transmission of encrypted credit card information, this attack could constitute the best mechanism ever invented for truly large-scale credit card fraud.

In fact, it is no exaggeration to say that if a determined criminal were more interested in vandalism than theft, he might be able to use this approach to destroy the entire global credit card infrastructure. (By a vandal, we mean someone who does not seek profit from his crime, but seeks instead to do the maximum amount of damage. A terrorist might well take this approach, for example.) Once such a criminal had obtained millions of valid credit card numbers, he could launch millions of transactions over the Internet in a single day, one on each of the stolen credit cards. It is difficult to be certain how well the existing credit card system would handle such a near-instantaneous increase of several orders of magnitude in the rate of credit card fraud, but it seems plausible that this could eliminate all profitability from the credit card industry, virtually overnight.

In short, while there are holes in the First Virtual system, they are small holes with bounded consequences, much like the risk in the physical world that is posed by people stealing credit card numbers from carbon papers in trash bins. Schemes that transmit encrypted credit card numbers trade those small risks for a much more catastrophic one. We think it's a very poor trade-off.

27 What are the real risks to worry about in Internet Commerce?

Given that every commerce mechanism that has ever been invented has had flaws, how should people assess the risks involved in Internet commerce? We believe that there are three key factors that need to be considered: the frequency with which the risk can be expected to turn into real examples of abuse; the level of inconvenience that users of the system experience when such abuse occurs; and the worst-case magnitude of the real effects of the abuse.

Only time and experience can really tell how frequently a system will actually be abused. First Virtual has been operating its system for a full year of staggering growth, with a fraud rate so low as to be almost completely negligible. In what is by far the most common abuse scenario -- where someone steals a Virtual PIN but doesn't intercept the email messages from First Virtual, so that the user promptly calls "fraud" -- the inconvenience to the user is small (a new account application is currently needed, though even this inconvenience can be reduced eventually) and the effects of the abuse are almost zero. The absolute worst-case abuse that we can imagine in our system would be if a criminal managed to break through all our security and steal our database of credit card numbers.

In contrast, what is a worst-case risk for First Virtual's server applies to EVERY MERCHANT'S MACHINE in some of the more naive cryptographic approaches, such as those that use SSL or something similar. In any scheme where credit cards are transmitted over the net and decoded on the merchant's machine, each merchant becomes an attractive target for criminals in a way that only First Virtual's most secure machine (which isn't even on the Internet) is a target in our system.

A few more sophisticated cryptographic schemes have reduced this vulnerability by moving the decryption to the acquirer's Internet-connected machines, which provides a small improvement by reducing the number of target machines for criminals. In this case, however, each acquiring bank needs to have a very high level of Internet security expertise, in order to achieve a level of safety commensurate with their level of liability. The early experience of banks on the Internet does little to inspire confidence that they will have that expertise any time soon.

Unfortunately, none of this is nearly the riskiest aspect of the cryptographic approaches. To return to our three key risk assessment factors: while only time and experience can answer the "frequency of abuse" question, the early experience with cryptography is not encouraging. There has been a steady stream of security holes and emergency fixes, of unexpected bugs and quick-fix patches. How many times can a cryptography vendor announce that it has fixed "the last major bug" and still retain credibility? Moreover, when cryptographic systems are compromised, users will generally have their underlying credit card rendered invalid due to fraud, a significantly greater annoyance than merely having a Virtual PIN rendered invalid. Finally, and most crucially, cryptographic systems can experience worst-case failures with absolutely catastrophic consequences, conceivably including the collapse of the entire credit card system, as describe in the previous section.

Credit card numbers were designed for easy machine recognition (they even include a checkdigit!) in an era when computers were rare and expensive, and only the good guys had them. Encryption seeks to prolong that era by making it difficult or expensive to intercept or recognize the credit card numbers, but in terms of Internet and computing technology, this is fundamentally a case of trying to swim upstream. The Virtual PIN mechanism was designed for the Internet era, where computation is cheap, communication is easily intercepted, viruses are easily deployed, and news of every software bug spreads like wildfire through the community of good citizens and would-be criminals alike. It should be no surprise that Virtual PINs provide the better risk management mechanism.

28 What is the real key to making Internet Commerce as safe as possible?

At the risk of being redundant, we must stress that nothing in life is completely safe. However, we can learn lessons from aerospace, medicine, and other technologies that are applied in high-risk situations. There are at least two fundamental risk management principles that we think should always be applied.

The first rule is to create redundancies and multiple backups. If a system has a single component which can cause a catastrophe if it fails, then that system will inevitably experience a catastrophic failure. To avoid this, in hardware systems, we build redundant physical systems, where the failure of one component activates a backup system in its place. In software, different strategies apply, because multiple copies of the same software would fail identically. The Space Shuttle, for example, has a backup software system that was written to the same specifications but coded entirely independently.

In making an Internet service safe, the key is to place multiple barriers in front of a would-be criminal. For example, in seeking the strongest security against Internet break-ins, a site generally builds multiple layers of firewalls, and sets various monitoring and alarm mechanisms after each one. Thus, if someone breaks through your first firewall, chances are good that they'll set off some alarms while trying to figure out how to break through the second one (assuming that it isn't an identical firewall, which would be virtually useless.)

First Virtual told you, in one of the questions above, precisely how to break the security of our transactions. What is most notable about this procedure, however, is that you have to do several different things in order to break security, the combination of which may expose you enough that you will be caught or traced. In contrast, schemes that encrypt credit cards and transmit them over the net have a single point of catastrophic failure. A criminal who finds a single bug in the encryption code, or who releases a virus that intercepts information as it is being typed in to the "secure" program, could collect credit card numbers in quantities that boggle the mind. Today's cryptographic systems do not, in general, have the kind of redundant and backup systems that engineers generally consider a requirement for truly critical applications. Most critically, the flaws in the cryptographic systems are far better suited to criminals making a large-scale automated approach to crime, whereas breaking First Virtual's system is far harder to automate, and thus pose a much smaller aggregated risk.

It's worth noting that the multi-step nature of breaking First Virtual's security is quite similar to the security applied to the physical currency produced by the US Department of the Treasury. It isn't enough to make a good forgery, for example, or even to steal the printing plates, because US money is printed on very special paper, the supply of which is tightly controlled. Thus, it takes multiple difficult actions to do a good job of forging currency, and the vulnerability of the overall system to any single point of failure (such as intercepting the paper supply or stealing the printing plates) is less than total.

The second rule is to keep things as simple as possible, especially for end users. Aerospace engineers now devote enormous effort to simplifying the displays and controls that pilots have to master. The average consumer is not likely to learn or follow any elaborate set of procedures or principles in his use of a payment system. Thus, over the years, the credit card companies have strived to put forth a set of very simple rules for the safe use of credit cards in various situations, such as "never give your credit card over the phone unless YOU placed the call."

First Virtual proposes an extremely simple rule for consumers: Never type your credit card, or any other sensitive financial information, into any Internet-connected machine. By contrast, advocates of cryptographic approaches find themselves telling consumers that it's OK to type such information into computers, but only if they're using trusted software. Unfortunately, it is impossible for most consumers to know which software can be trusted, and which software has been silently compromised by an invisible virus, modified by a malicious system administrator, or simply sold by the vendor with an accidental but fatal software error. There's simply no way to tell naive users of cryptographic systems how to use those systems safely.

29 What are the costs, risks, and benefits involved in using the First Virtual system?

As stated previously, we believe that advocates of a given system for Internet commerce should be clear, explicit, and up-front about the costs and risks of using their system, as well as the benefits. We try to be very straightforward about the costs and risks involved in our system, and we encourage merchants and consumers alike to demand a similar level of "full disclosure" from other systems.

The costs of using First Virtual are straightforward on an economic level: There is a $10 fee to set up a seller's account, or a $2 fee for a buyer's account. (The fee recurs if the underlying financial information changes.) Beyond that, the seller pays 29 cents plus 2% of each transaction that is processed for him by First Virtual, and pays a $1 fee each time an aggregated deposit is made to his bank account.

The less direct costs are also easily detailed: The 91-day fund holding period for merchants who have not qualified for an exemption is clearly a serious cost. Non-economic costs involved in using the system include the lack of an automatic direct identification of the consumer. (First Virtual transactions are not anonymous in that First Virtual can trace them readily, but a Virtual PIN neatly shields the buyer's identity from the seller if that is the buyer's wish. Of course, a seller can always ask for the buyer to

identify himself.) Another cost that must be considered is whatever effort is involved in setting up one's system to take First Virtual, although this cost exists in any Internet commerce system.

The risks involved in using the First Virtual system are few, but important to understand. Those merchants who sell information in the manner First Virtual recommends -- delivering the information before the buyer says yes -- risk occasional non-payment for information goods. The use of stolen credit cards to set up a First Virtual account is a risk, but one that is certainly no greater than for mail order or telephone order businesses. Virtual PINs are specifically designed to pose low risk in the event of their theft, unlike credit cards, but a stolen Virtual PIN can be used to obtain information from a merchant who does not wait for the buyer's "yes" reply. Finally, like all approaches to Internet commerce, the system is vulnerable to attack from a criminal who manages to obtain system administration privileges on a critical machine, notably one where several First Virtual customers receive their mail. The risk in that case, for each Virtual PIN compromised, is roughly equivalent to the theft of a single credit card, but only for use on the Internet. Most important, there are no known risks in the system that leave it vulnerable to large-scale automated fraud.

The benefits of First Virtual's system are many. Payment may be accepted over the Internet from buyers who have no special software. The system's simple, open protocol specification has quickly led to the development of third-party merchant tools, which gives merchants more options and more freedom. Merchants never have to deal with actual credit card processing, but instead have money deposited directly into their bank accounts by First Virtual. Most unusually, there is no screening process for becoming a First Virtual seller -- the system is entrepreneur-friendly.

Finally, one of the greatest benefits of the First Virtual system is that it does not suffer from certain risks or costs that are common to most other proposed systems for Internet commerce. It does not make merchants into attractive targets for criminals by making them collect credit card numbers. It does not make buyers obtain, install, and learn to use new software, and therefore there is no need for merchants or banks to worry about answering users' questions about such software. It does not impose a qualification process on would-be entrepreneurs. Because it does not use encryption, it does not risk catastrophic failure in the event that a single algorithm or program proves to be flawed. There are no export or licensing restrictions involved in the use of the system. There is no need to worry about the ability of consumers to safely protect and manage secret cryptographic keys. There is no scenario in which fraud can be automated, or in which a flaw in the system can lead to massive losses for credit card issuing or acquiring banks. Encryption is a radically new technology for the consumer market, and providers of financial services should be at least a bit conservative. By avoiding the cryptographic bells and whistles with which most other approaches are weighted down, the First Virtual system also avoids an enormous range of poorly-understood costs and risks that those systems have yet to squarely face.

In short, First Virtual is a simple, safe, and conservative approach to Internet commerce. The biggest problem faced in its operation has been the management of its exponential growth. After a full year of operation, no criminal has broken it, no fatal flaws have been discovered, and no headlines have been made about unexpected security holes. The system works simply, so it simply works.

Postscript

If you've read this far, you can tell that we've spent a lot of time thinking about how to make Internet transactions safe and simple. We don't expect everyone to agree with every point we've made, and we're very interested in any feedback you might have. We're especially interested in how to make our system even safer than it is today. If you want to discuss any of the issues raised in this document, we encourage you to join our "fv-users" mailing list, where we enjoy a continuing free-wheeling and uncensored discussion. You can join the list by sending mail to "fv-users-request@fv.com" with the word "subscribe" as the subject of your mail. We would welcome any input you cared to provide.