

Perils and Pitfalls of Practical CyberCommerce

The Lessons of First Virtual's First Year

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

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

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 use any procedure or protocol to meet and decide to transact business. 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 a full year of operating that system, processing transactions for real money. However, the system is sufficiently different from most other proposed approaches to Internet commerce that we have prepared a rather lengthy list of commonly-raised concerns about First Virtual, and our responses, as Appendix A to an extended version of this paper, available on our web pages at <http://www.fv.com/pubdocs/fv-austin.txt>.

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.

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 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 in fact take place.)

More generally, 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 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 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 distributed 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 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.

²Over 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.

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."

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 some of the most well-known software vendors.

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 currently 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. 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.

Software configuration bugs: The widely-used Netscape browser, for example, can be used to send mail, but, in its configuration-setting mechanism, makes no attempt to verify that the user-supplied email address is correct (or even syntactically legal!) Thus a surprising number of Netscape users never receive any replies to their email, and never know why.

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 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. 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. Not surprisingly, we have given these issues a great deal of thought in recent months, and have reached some tentative conclusions.

The major risk in cryptography is the compromise of the cryptographic keys. 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. Although 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, our operational experience leaves us skeptical that 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. The longer-lived the keys are, the more likely it is that an attack will undermine their value. This is an area with crucial security consequences which are often 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 lifetime is probably considerably more vulnerable to criminal attack than a 512-bit key with a 1-month lifetime.

In assessing the importance of the various risks, it is important to distinguish between the two main applications of cryptographic technology: authentication (digital signatures) 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.

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.

Technically, 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". A realistic analysis of any cryptographic commerce mechanism must 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, stealing the secret keys involved, or finding a serious bug in a widely-used software 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 relatively common.

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. 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 possibly 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.

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. 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.

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.

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.

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.