

Network Working Group
OBSOLETE [RFC 1750](#)

Expires April 2005

Donald E. Eastlake, 3rd
Jeffrey I. Schiller
Steve Crocker
October 2004

Randomness Requirements for Security

<[draft-eastlake-randomness2-09.txt](#)>

Status of This Document

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, or will be disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

This document is intended to become a Best Current Practice. Comments should be sent to the authors. Distribution is unlimited.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than a "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Copyright (C) The Internet Society 2004. All Rights Reserved.

Abstract

Security systems are built on strong cryptographic algorithms that foil pattern analysis attempts. However, the security of these systems is dependent on generating secret quantities for passwords, cryptographic keys, and similar quantities. The use of pseudo-random processes to generate secret quantities can result in pseudo-security. The sophisticated attacker of these security systems may find it easier to reproduce the environment that produced the secret quantities, searching the resulting small set of possibilities, than

to locate the quantities in the whole of the potential number space.
Choosing random quantities to foil a resourceful and motivated

adversary is surprisingly difficult. This document points out many pitfalls in using traditional pseudo-random number generation techniques for choosing such quantities. It recommends the use of truly random hardware techniques and shows that the existing hardware on many systems can be used for this purpose. It provides suggestions to ameliorate the problem when a hardware solution is not available. And it gives examples of how large such quantities need to be for some applications.

Acknowledgements

Special thanks to Paul Hoffman and John Kelsey for their extensive comments and to Peter Gutmann, who has permitted the incorporation of material from his paper "Software Generation of Practically Strong Random Numbers".

The following other persons (in alphabetic order) have also contributed substantially to this document:

Daniel Brown, Don Davis, Peter Gutmann, Tony Hansen, Sandy Harris, Paul Hoffman, Scott Hollenback, Russ Housley, Christian Huitema, John Kelsey, and Damir Rajnovic.

The following persons (in alphabetic order) contributed to [RFC 1750](#), the predecessor of this document:

David M. Balenson, Don T. Davis, Carl Ellison, Marc Horowitz, Christian Huitema, Charlie Kaufman, Steve Kent, Hal Murray, Neil Haller, Richard Pitkin, Tim Redmond, and Doug Tygar.

INTERNET DRAFT Randomness Requirements for Security October 2004

Table of Contents

Status of This Document.....	1
Abstract.....	1
Acknowledgements.....	2
Table of Contents.....	3
1 . Introduction.....	5
2 . General Requirements.....	6
3 . Traditional Pseudo-Random Sequences.....	9
4 . Unpredictability.....	11
4.1 Problems with Clocks and Serial Numbers.....	11
4.2 Timing and Value of External Events.....	12
4.3 The Fallacy of Complex Manipulation.....	12
4.4 The Fallacy of Selection from a Large Database.....	13
5 . Hardware for Randomness.....	15
5.1 Volume Required.....	15
5.2 Sensitivity to Skew.....	15
5.2.1 Using Stream Parity to De-Skew.....	16
5.2.2 Using Transition Mappings to De-Skew.....	17
5.2.3 Using FFT to De-Skew.....	18
5.2.4 Using Compression to De-Skew.....	18
5.3 Existing Hardware Can Be Used For Randomness.....	19
5.3.1 Using Existing Sound/Video Input.....	19
5.3.2 Using Existing Disk Drives.....	19
5.4 Ring Oscillator Sources.....	20

6. Recommended Software Strategy.....	22
6.1 Mixing Functions.....	22
6.1.1 A Trivial Mixing Function.....	22
6.1.2 Stronger Mixing Functions.....	23
6.1.3 Using S-Boxes for Mixing.....	25
6.1.4 Diffie-Hellman as a Mixing Function.....	25
6.1.5 Using a Mixing Function to Stretch Random Bits.....	25
6.1.6 Other Factors in Choosing a Mixing Function.....	26
6.2 Non-Hardware Sources of Randomness.....	27
6.3 Cryptographically Strong Sequences.....	28
6.3.1 OFB and CTR Sequences.....	28
6.3.2 The Blum Blum Shub Sequence Generator.....	29
6.3.3 Entropy Pool Techniques.....	30
7. Key Generation Examples and Standards.....	32
7.1 US DoD Recommendations for Password Generation.....	32
7.2 X9.17 Key Generation.....	32
7.3 DSS Pseudo-Random Number Generation.....	33

D. Eastlake, J. Schiller, S. Crocker

[Page 3]

INTERNET DRAFT Randomness Requirements for Security October 2004

7.4 X9.82 Pseudo-Random Number Generation.....	34
7.5 The /dev/random Device.....	34
7.6 Windows CryptGenRandom.....	36
8. Examples of Randomness Required.....	37
8.1 Password Generation.....	37
8.2 A Very High Security Cryptographic Key.....	38
8.2.1 Effort per Key Trial.....	38
8.2.2 Meet in the Middle Attacks.....	39
8.2.3 Other Considerations.....	40
9. Conclusion.....	41
10. Security Considerations.....	42
11. Copyright and Disclaimer.....	42
12. Appendix A: Changes from RFC 1750.....	43
14. Informative References.....	44
Author's Addresses.....	48
File Name and Expiration.....	48

1. Introduction

Software cryptography is coming into wider use and is continuing to spread, although there is a long way to go until it becomes pervasive.

Systems like SSH, IPSEC, TLS, S/MIME, PGP, DNSSEC, Kerberos, etc. are maturing and becoming a part of the network landscape [SSH, IPSEC, MAIL*, TLS, DNSSEC]. By comparison, when the previous version of this document [[RFC 1750](#)] was issued in 1994, about the only Internet cryptographic security specification in the IETF was the Privacy Enhanced Mail protocol [MAIL PEM *].

These systems provide substantial protection against snooping and spoofing. However, there is a potential flaw. At the heart of all

cryptographic systems is the generation of secret, unguessable (i.e., random) numbers.

The lack of generally available facilities for generating such random numbers, that is the lack of general availability of truly unpredictable sources, forms an open wound in the design of cryptographic software. For the software developer who wants to build a key or password generation procedure that runs on a wide range of hardware, this is a very real problem.

It is important to keep in mind that the requirement is for data that an adversary has a very low probability of guessing or determining. This can easily fail if pseudo-random data is used which only meets traditional statistical tests for randomness or which is based on limited range sources, such as clocks. Sometimes such pseudo-random quantities are determinable by an adversary searching through an embarrassingly small space of possibilities.

This Best Current Practice describes techniques for producing random quantities that will be resistant to such attack. It recommends that future systems include hardware random number generation or provide access to existing hardware that can be used for this purpose. It suggests methods for use if such hardware is not available. And it gives some estimates of the number of random bits required for sample applications.

[2. General Requirements](#)

A commonly encountered randomness requirement today is the user password. This is usually a simple character string. Obviously, if a password can be guessed, it does not provide security. (For re-usable passwords, it is desirable that users be able to remember the password. This may make it advisable to use pronounceable character

strings or phrases composed on ordinary words. But this only affects the format of the password information, not the requirement that the password be very hard to guess.)

Many other requirements come from the cryptographic arena. Cryptographic techniques can be used to provide a variety of services including confidentiality and authentication. Such services are based on quantities, traditionally called "keys", that are unknown to and unguessable by an adversary.

Generally speaking, the above two examples also illustrate two different types of random quantities that may be wanted. In the case of human usable passwords, the only important characteristic is that it be unguessable; it is not important that they may be composed of ASCII characters, for example, so the top bit of every byte is zero. On the other hand, for fixed length keys and the like, you normally want quantities that are indistinguishable from truly random, that is, all bits will pass statistical randomness tests.

In some cases, such as the use of symmetric encryption with the one time pads or an algorithm like the US Advanced Encryption Standard [\[AES\]](#), the parties who wish to communicate confidentially and/or with authentication must all know the same secret key. In other cases, using what are called asymmetric or "public key" cryptographic techniques, keys come in pairs. One key of the pair is private and must be kept secret by one party, the other is public and can be published to the world. It is computationally infeasible to determine the private key from the public key and knowledge of the public is of no help to an adversary [\[ASYMMETRIC\]](#). [\[SCHNEIER, FERGUSON, KAUFMAN\]](#)

The frequency and volume of the requirement for random quantities differs greatly for different cryptographic systems. Using pure RSA, random quantities are required only when a new key pair is generated; thereafter any number of messages can be signed without a further need for randomness. The public key Digital Signature Algorithm devised by the US National Institute of Standards and Technology (NIST) requires good random numbers for each signature [\[DSS\]](#). And encrypting with a one time pad, in principle the strongest possible encryption technique, requires a volume of randomness equal to all the messages to be processed. [\[SCHNEIER, FERGUSON, KAUFMAN\]](#)

In most of these cases, an adversary can try to determine the "secret" key by trial and error. (This is possible as long as the key

is enough smaller than the message that the correct key can be uniquely identified.) The probability of an adversary succeeding at this must be made acceptably low, depending on the particular application. The size of the space the adversary must search is related to the amount of key "information" present in the information theoretic sense [[SHANNON](#)]. This depends on the number of different secret values possible and the probability of each value as follows:

$$\text{Bits-of-information} = \frac{\sum_{i=1}^n -p_i \log_2(p_i)}{1}$$

where i counts from 1 to the number of possible secret values and $p_{\text{sub } i}$ is the probability of the value numbered i . (Since $p_{\text{sub } i}$ is less than one, the log will be negative so each term in the sum will be non-negative.)

If there are 2^n different values of equal probability, then n bits of information are present and an adversary would, on the average, have to try half of the values, or $2^{(n-1)}$, before guessing the secret quantity. If the probability of different values is unequal, then there is less information present and fewer guesses will, on average, be required by an adversary. In particular, any values that the adversary can know are impossible, or are of low probability, can be initially ignored by an adversary, who will search through the more probable values first.

For example, consider a cryptographic system that uses 128 bit keys. If these 128 bit keys are derived by using a fixed pseudo-random number generator that is seeded with an 8 bit seed, then an adversary needs to search through only 256 keys (by running the pseudo-random number generator with every possible seed), not the 2^{128} keys that may at first appear to be the case. Only 8 bits of "information" are in these 128 bit keys.

While the above analysis is correct on average, it can be misleading in some cases for cryptographic analysis where what is really important is the work factor for an adversary. For example, assume that there was a pseudo-random number generator generating 128 bit keys, as in the previous paragraph, but that it generated 0 half of the time and a random selection from the remaining $2^{128} - 1$ values the rest of the time. The Shannon equation above says that there are 64 bits of information in one of these key values but an adversary, by simply trying the values 0, can break the security of half of the uses, albeit a random half. Thus for cryptographic purposes, it is also useful to look at other measures, such as min-entropy, defined as

INTERNET DRAFT

Randomness Requirements for Security

October 2004

$$\text{Min-entropy} = -\log \left(\max_i (p_i) \right)$$

where i is as above. Using this equation, we get 1 bit of min-entropy for our new hypothetical distribution as opposed to 64 bits of classical Shannon entropy.

A continuous spectrum of entropies, sometimes called Renyi entropy, have been defined, specified by a parameter r . When $r = 1$, it is Shannon entropy, and with $r = \text{infinity}$, it is min-entropy. When $r = 0$, it is just $\log(n)$ where n is the number of non-zero probabilities. Renyi entropy is a non-increasing function of r , so min-entropy is always the most conservative measure of entropy and usually the best to use for cryptographic evaluation. [[LUBY](#)]

3. Traditional Pseudo-Random Sequences

This section talks about traditional sources of deterministic of "pseudo-random" numbers. These typically start with a "seed" quantity and use numeric or logical operations to produce a sequence of values. Note that none of the techniques discussed in this section is suitable for cryptographic use. They are presented for general information.

[KNUTH] has a classic exposition on pseudo-random numbers. Applications he mentions are simulation of natural phenomena, sampling, numerical analysis, testing computer programs, decision making, and games. None of these have the same characteristics as the sort of security uses we are talking about. Only in the last two could there be an adversary trying to find the random quantity. However, in these cases, the adversary normally has only a single chance to use a guessed value. In guessing passwords or attempting to break an encryption scheme, the adversary normally has many, perhaps unlimited, chances at guessing the correct value. They can store the message they are trying to break and repeatedly attack it. They are also be assumed to be aided by a computer.

For testing the "randomness" of numbers, Knuth suggests a variety of measures including statistical and spectral. These tests check things like autocorrelation between different parts of a "random" sequence or distribution of its values. But they could be met by a constant stored random sequence, such as the "random" sequence printed in the CRC Standard Mathematical Tables [[CRC](#)]. Despite meeting all the tests suggested by Knuth, that sequence is unsuitable for cryptographic use as adversaries must be assumed to have copies of all common published "random" sequences and will be able to spot the source and predict future values.

A typical pseudo-random number generation technique, known as a

linear congruence pseudo-random number generator, is modular arithmetic where the value numbered N+1 is calculated from the value numbered N by

$$V_{N+1} = (V_N * a + b) \text{ (Mod } c)$$

The above technique has a strong relationship to linear shift register pseudo-random number generators, which are well understood cryptographically [SHIFT*]. In such generators bits are introduced at one end of a shift register as the Exclusive Or (binary sum without carry) of bits from selected fixed taps into the register. For example:


D. Eastlake, J. Schiller, S. Crocker

[Page 9]

INTERNET DRAFT

Randomness Requirements for Security

October 2004



$$V_{N+1} = ((V_N * 2) + B_0 \text{ .xor. } B_2 \dots) \text{ (Mod } 2^n)$$

The goodness of traditional pseudo-random number generator algorithms is measured by statistical tests on such sequences. Carefully chosen values a , b , c , and initial V or the placement of shift register tap in the above simple processes can produce excellent statistics.

These sequences may be adequate in simulations (Monte Carlo experiments) as long as the sequence is orthogonal to the structure of the space being explored. Even there, subtle patterns may cause problems. However, such sequences are clearly bad for use in security applications. They are fully predictable if the initial state is known. Depending on the form of the pseudo-random number generator,

the sequence may be determinable from observation of a short portion of the sequence [[SCHNEIER](#), [STERN](#)]. For example, with the generators above, one can determine $V(n+1)$ given knowledge of $V(n)$. In fact, it has been shown that with these techniques, even if only one bit of the pseudo-random values are released, the seed can be determined from short sequences.

Not only have linear congruent generators been broken, but techniques are now known for breaking all polynomial congruent generators. [[KRAWCZYK](#)]

[4](#). Unpredictability

Statistically tested randomness in the traditional sense described in [section 3](#) is NOT the same as the unpredictability required for security use.

For example, use of a widely available constant sequence, such as that from the CRC tables, is very weak against an adversary. Once they learn of or guess it, they can easily break all security, future and past, based on the sequence. [[CRC](#)] Yet the statistical properties of these tables are good. So you should keep in mind that passing statistical tests doesn't tell you that something is unpredictable.

The following sections describe the limitations of some randomness generation techniques and sources. Much better sources are described in [Section 5](#).

[4.1](#) Problems with Clocks and Serial Numbers

Computer clocks, or similar operating system or hardware values, provide significantly fewer real bits of unpredictability than might appear from their specifications.

Tests have been done on clocks on numerous systems and it was found that their behavior can vary widely and in unexpected ways. One version of an operating system running on one set of hardware may actually provide, say, microsecond resolution in a clock while a different configuration of the "same" system may always provide the same lower bits and only count in the upper bits at much lower resolution. This means that successive reads on the clock may produce identical values even if enough time has passed that the value "should" change based on the nominal clock resolution. There are also cases where frequently reading a clock can produce artificial sequential values because of extra code that checks for the clock being unchanged between two reads and increases it by one! Designing portable application code to generate unpredictable numbers based on such system clocks is particularly challenging because the system designer does not always know the properties of the system clocks that the code will execute on.

Use of hardware serial numbers such as an Ethernet addresses may also provide fewer bits of uniqueness than one would guess. Such quantities are usually heavily structured and subfields may have only a limited range of possible values or values easily guessable based on approximate date of manufacture or other data. For example, it is likely that a company that manufactures both computers and Ethernet adapters will, at least internally, use its own adapters, which significantly limits the range of built-in addresses.

Problems such as those described above related to clocks and serial numbers make code to produce unpredictable quantities difficult if the code is to be ported across a variety of computer platforms and systems.

[4.2](#) Timing and Value of External Events

It is possible to measure the timing and content of mouse movement,

key strokes, and similar user events. This is a reasonable source of unguessable data with some qualifications. On some machines, inputs such as key strokes are buffered. Even though the user's inter-keystroke timing may have sufficient variation and unpredictability, there might not be an easy way to access that variation. Another problem is that no standard method exists to sample timing details. This makes it hard to build standard software intended for distribution to a large range of machines based on this technique.

The amount of mouse movement or the keys actually hit are usually easier to access than timings but may yield less unpredictability as the user may provide highly repetitive input.

Other external events, such as network packet arrival times and lengths, can also be used, but only with great care. In particular, the possibility of manipulation of such network traffic measurements by an adversary and the lack of history at system start up must be carefully considered. If this input is subject to manipulation, it must not be trusted as a source of entropy.

Indeed, almost any external sensor, such as raw radio reception or temperature sensing in appropriately equipped computers, can be used in principle. But in each case careful consideration must be given to how much such data is subject to adversarial manipulation and to how much entropy it can actually provide.

The above techniques are quite powerful against any attackers having no access to the quantities being measured. For example, they would be powerful against offline attackers who had no access to your environment and were trying to crack your random seed after the fact. In all cases, the more accurately you can measure the timing or value of an external sensor, the more rapidly you can generate bits.

[4.3](#) The Fallacy of Complex Manipulation

One strategy which may give a misleading appearance of unpredictability is to take a very complex algorithm (or an excellent traditional pseudo-random number generator with good statistical

properties) and calculate a cryptographic key by starting with limited data such as the computer system clock value as the seed. An

adversary who knew roughly when the generator was started would have a relatively small number of seed values to test as they would know likely values of the system clock. Large numbers of pseudo-random bits could be generated but the search space an adversary would need to check could be quite small.

Thus very strong and/or complex manipulation of data will not help if the adversary can learn what the manipulation is and there is not enough unpredictability in the starting seed value. They can usually use the limited number of results stemming from a limited number of seed values to defeat security.

Another serious strategy error is to assume that a very complex pseudo-random number generation algorithm will produce strong random numbers when there has been no theory behind or analysis of the algorithm. There is an excellent example of this fallacy right near the beginning of Chapter 3 in [\[KNUTH\]](#) where the author describes a complex algorithm. It was intended that the machine language program corresponding to the algorithm would be so complicated that a person trying to read the code without comments wouldn't know what the program was doing. Unfortunately, actual use of this algorithm showed that it almost immediately converged to a single repeated value in one case and a small cycle of values in another case.

Not only does complex manipulation not help you if you have a limited range of seeds but blindly chosen complex manipulation can destroy the randomness in a good seed!

[4.4](#) The Fallacy of Selection from a Large Database

Another strategy that can give a misleading appearance of unpredictability is selection of a quantity randomly from a database and assume that its strength is related to the total number of bits in the database. For example, typical USENET servers process many megabytes of information per day [\[USENET\]](#). Assume a random quantity was selected by fetching 32 bytes of data from a random starting point in this data. This does not yield $32 \times 8 = 256$ bits worth of unguessability. Even after allowing that much of the data is human language and probably has no more than 2 or 3 bits of information per byte, it doesn't yield $32 \times 2 = 64$ bits of unguessability. For an adversary with access to the same usenet database the unguessability rests only on the starting point of the selection. That is perhaps a little over a couple of dozen bits of unguessability.

The same argument applies to selecting sequences from the data on a publicly available CD/DVD recording or any other large public

database. If the adversary has access to the same database, this "selection from a large volume of data" step buys little. However, if a selection can be made from data to which the adversary has no access, such as system buffers on an active multi-user system, it may be of help.

[5](#). Hardware for Randomness

Is there any hope for true strong portable randomness in the future? There might be. All that's needed is a physical source of unpredictable numbers.

A thermal noise (sometimes called Johnson noise in integrated circuits) or radioactive decay source and a fast, free-running oscillator would do the trick directly [[GIFFORD](#)]. This is a trivial amount of hardware, and could easily be included as a standard part of a computer system's architecture. Most audio (or video) input devices are useable [[TURBID](#)]. Furthermore, any system with a spinning disk or ring oscillator and a stable (crystal) time source or the like has an adequate source of randomness ([[DAVIS](#)] and [Section 5.4](#)). All that's needed is the common perception among computer vendors that this small additional hardware and the software to access it is necessary and useful.

[5.1](#) Volume Required

How much unpredictability is needed? Is it possible to quantify the requirement in, say, number of random bits per second?

The answer is not very much is needed. For AES, the key can be 128 bits and, as we show in an example in [Section 8](#), even the highest security system is unlikely to require strong keying material of much over 200 bits. If a series of keys are needed, they can be generated from a strong random seed (starting value) using a cryptographically strong sequence as explained in [Section 6.3](#). A few hundred random bits generated at start up or once a day would be enough using such techniques. Even if the random bits are generated as slowly as one per second and it is not possible to overlap the generation process, it should be tolerable in most high security applications to wait 200 seconds occasionally.

These numbers are trivial to achieve. It could be done by a person

repeatedly tossing a coin. Almost any hardware based process is likely to be much faster.

[5.2](#) Sensitivity to Skew

Is there any specific requirement on the shape of the distribution of the random numbers? The good news is the distribution need not be uniform. All that is needed is a conservative estimate of how non-uniform it is to bound performance. Simple techniques to de-skew the bit stream are given below and stronger cryptographic techniques are

D. Eastlake, J. Schiller, S. Crocker

[Page 15]

INTERNET DRAFT

Randomness Requirements for Security

October 2004

described in [Section 6.1.2](#) below.

[5.2.1](#) Using Stream Parity to De-Skew

As a simple but not particularly practical example, consider taking a sufficiently long string of bits and map the string to "zero" or "one". The mapping will not yield a perfectly uniform distribution, but it can be as close as desired. One mapping that serves the purpose is to take the parity of the string. This has the advantages that it is robust across all degrees of skew up to the estimated maximum skew and is absolutely trivial to implement in hardware.

The following analysis gives the number of bits that must be sampled:

Suppose the ratio of ones to zeros is $(0.5 + e)$ to $(0.5 - e)$, where e is between 0 and 0.5 and is a measure of the "eccentricity" of the distribution. Consider the distribution of the parity function of N bit samples. The probabilities that the parity will be one or zero will be the sum of the odd or even terms in the binomial expansion of $(p + q)^N$, where $p = 0.5 + e$, the probability of a one, and $q = 0.5 - e$, the probability of a zero.

These sums can be computed easily as

$$\text{and } 1/2 * \left(\sum_{N=0}^N \binom{N}{k} p^k q^{N-k} + \sum_{N=0}^N \binom{N}{k} p^k q^{N-k} \right)$$

$$1/2 * ((p + q)^N - (p - q)^N).$$

(Which one corresponds to the probability the parity will be 1 depends on whether N is odd or even.)

Since $p + q = 1$ and $p - q = 2e$, these expressions reduce to

$$1/2 * [1 + (2e)^N]$$

and

$$1/2 * [1 - (2e)^N].$$

Neither of these will ever be exactly 0.5 unless e is zero, but we can bring them arbitrarily close to 0.5. If we want the probabilities to be within some delta d of 0.5, i.e. then

$$(0.5 + (0.5 * (2e)^N)) < 0.5 + d.$$

Solving for N yields $N > \log(2d)/\log(2e)$. (Note that 2e is less than 1, so its log is negative. Division by a negative number reverses the sense of an inequality.)

The following table gives the length of the string which must be sampled for various degrees of skew in order to come within 0.001 of a 50/50 distribution.

Prob(1)	e	N
0.5	0.00	1
0.6	0.10	4
0.7	0.20	7
0.8	0.30	13
0.9	0.40	28
0.95	0.45	59
0.99	0.49	308

The last entry shows that even if the distribution is skewed 99% in favor of ones, the parity of a string of 308 samples will be within

0.001 of a 50/50 distribution. But, as we shall see in [section 6.1.2](#), there are much stronger techniques that extract more of the available entropy.

[5.2.2](#) Using Transition Mappings to De-Skew

Another technique, originally due to von Neumann [VON NEUMANN], is to examine a bit stream as a sequence of non-overlapping pairs. You could then discard any 00 or 11 pairs found, interpret 01 as a 0 and 10 as a 1. Assume the probability of a 1 is $0.5+e$ and the probability of a 0 is $0.5-e$ where e is the eccentricity of the source and described in the previous section. Then the probability of each pair is as follows:

pair		probability	
00	$(0.5 - e)^2$	$= 0.25 - e + e^2$	
01	$(0.5 - e) * (0.5 + e)$	$= 0.25 - e^2$	
10	$(0.5 + e) * (0.5 - e)$	$= 0.25 - e^2$	
11	$(0.5 + e)^2$	$= 0.25 + e + e^2$	

This technique will completely eliminate any bias but at the expense of taking an indeterminate number of input bits for any particular

desired number of output bits. The probability of any particular pair being discarded is $0.5 + 2e^2$ so the expected number of input bits to produce X output bits is $X/(0.25 - e^2)$.

This technique assumes that the bits are from a stream where each bit has the same probability of being a 0 or 1 as any other bit in the stream and that bits are not correlated, i.e., that the bits are identical independent distributions. If alternate bits were from two correlated sources, for example, the above analysis breaks down.

The above technique also provides another illustration of how a simple statistical analysis can mislead if one is not always on the lookout for patterns that could be exploited by an adversary. If the algorithm were mis-read slightly so that overlapping successive bits

pairs were used instead of non-overlapping pairs, the statistical analysis given is the same; however, instead of providing an unbiased uncorrelated series of random 1s and 0s, it instead produces a totally predictable sequence of exactly alternating 1s and 0s.

[5.2.3](#) Using FFT to De-Skew

When real world data consists of strongly biased or correlated bits, it may still contain useful amounts of randomness. This randomness can be extracted through use of various transforms, the most powerful of which are described in [section 6.1.2](#) below.

Using the Fourier transform of the data or its optimized variant, the FFT, is a technique interesting primarily for theoretical reasons. It can be shown that this will discard strong correlations. If adequate data is processed and remaining correlations decay, spectral lines approaching statistical independence and normally distributed randomness can be produced [[BRILLINGER](#)].

[5.2.4](#) Using Compression to De-Skew

Reversible compression techniques also provide a crude method of de-skewing a skewed bit stream. This follows directly from the definition of reversible compression and the formula in [Section 2](#) above for the amount of information in a sequence. Since the compression is reversible, the same amount of information must be present in the shorter output than was present in the longer input. By the Shannon information equation, this is only possible if, on average, the probabilities of the different shorter sequences are more uniformly distributed than were the probabilities of the longer sequences. Therefore the shorter sequences must be de-skewed relative to the input.

However, many compression techniques add a somewhat predictable preface to their output stream and may insert such a sequence again periodically in their output or otherwise introduce subtle patterns of their own. They should be considered only a rough technique compared with those described in [Section 6.1.2](#). At a minimum, the beginning of the compressed sequence should be skipped and only later

bits used for applications requiring roughly random bits.

[5.3](#) Existing Hardware Can Be Used For Randomness

As described below, many computers come with hardware that can, with care, be used to generate truly random quantities.

[5.3.1](#) Using Existing Sound/Video Input

Many computers are built with inputs that digitize some real world analog source, such as sound from a microphone or video input from a camera. Under appropriate circumstances, such input can provide reasonably high quality random bits. The "input" from a sound digitizer with no source plugged in or a camera with the lens cap on, if the system has enough gain to detect anything, is essentially thermal noise. This method is extremely hardware and implementation dependent.

For example, on some UNIX based systems, one can read from the `/dev/audio` device with nothing plugged into the microphone jack or the microphone receiving only low level background noise. Such data is essentially random noise although it should not be trusted without some checking in case of hardware failure. It will, in any case, need to be de-skewed as described elsewhere.

Combining this with compression to de-skew one can, in UNIXese, generate a huge amount of medium quality random data by doing

```
cat /dev/audio | compress - >random-bits-file
```

A detailed examination of this type of randomness source appears in [\[TURBID\]](#).

[5.3.2](#) Using Existing Disk Drives

Disk drives have small random fluctuations in their rotational speed due to chaotic air turbulence [\[DAVIS, Jakobsson\]](#). By adding low level disk seek time instrumentation to a system, a series of

measurements can be obtained that include this randomness. Such data is usually highly correlated so that significant processing is needed, such as described in 6.1.2 below. Nevertheless experimentation a decade ago showed that, with such processing, even slow disk drives on the slower computers of that day could easily produce 100 bits a minute or more of excellent random data.

Every increase in processor speed, which increases the resolution with which disk motion can be timed, or increase in the rate of disk seeks, increases the rate of random bit generation possible with this technique. At the time of this paper and using modern hardware, a more typical rate of random bit production would be in excess of 10,000 bits a second. This technique is used in many operating system library random number generators.

Note: the inclusion of cache memories in disk controllers has little effect on this technique if very short seek times, which represent cache hits, are simply ignored.

[5.4](#) Ring Oscillator Sources

If an integrated circuit is being designed or field programmed, an odd number of gates can be connected in series to produce a free-running ring oscillator. By sampling a point in the ring at a fixed frequency, say one determined by a stable crystal oscillator, some amount of entropy can be extracted due to variations in the free-running oscillator timing. It is possible to increase the rate of entropy by xor'ing sampled values from a few ring oscillators with relatively prime lengths. It is sometimes recommended that an odd number of rings be used so that, even if the rings somehow become synchronously locked to each other, there will still be sampled bit transitions. Another possibility source to sample is the output of a noisy diode.

Sampled bits from such sources will have to be heavily de-skewed, as disk rotation timings must be ([Section 5.3.2](#)). An engineering study would be needed to determine the amount of entropy being produced depending on the particular design. In any case, these can be good sources whose cost is a trivial amount of hardware by modern standards.

As an example, IEEE 802.11i suggests that the circuit below be considered, with due attention in the design to isolation of the rings from each other and from clocked circuits to avoid undesired synchronization, etc., and extensive post processing. [IEEE 802.11i]

INTERNET DRAFT

Randomness Requirements for Security

October 2004

```

      | \      | \      | \
+-->| >0-->| >0-- 19 total --| >0--+-----+
|   | /      | /      | /      |         |
|   |         |         |         |         |
+-----+-----+-----+         V
                                     +-----+
      | \      | \      | \      |   |   |   |
+-->| >0-->| >0-- 23 total --| >0--+-->| XOR |-----> output
|   | /      | /      | /      |   |   |   |
|   |         |         |         |   |   |   |
+-----+-----+-----+         ^  ^
                                     |   |
      | \      | \      | \      |   |   |
+-->| >0-->| >0-- 29 total --| >0--+-----+   |
|   | /      | /      | /      |   |   |   |
|   |         |         |         |   |   |   |
+-----+-----+-----+   |   |   |
                                |   |
other randomness if available-----+

```


[6.](#) Recommended Software Strategy

What is the best overall strategy for meeting the requirement for unguessable random numbers in the absence of a reliable hardware source? It is to obtain random input from a number of uncorrelated sources and to mix them with a strong mixing function. Such a function will preserve the randomness present in any of the sources even if other quantities being combined happen to be fixed or easily guessable. This may be advisable even with a good hardware source, as hardware can also fail, though this should be weighed against any increase in the chance of overall failure due to added software complexity.

[6.1](#) Mixing Functions

A strong mixing function is one which combines inputs and produces an output where each output bit is a different complex non-linear function of all the input bits. On average, changing any input bit will change about half the output bits. But because the relationship is complex and non-linear, no particular output bit is guaranteed to change when any particular input bit is changed.

Consider the problem of converting a stream of bits that is skewed towards 0 or 1 or which has a somewhat predictable pattern to a shorter stream which is more random, as discussed in [Section 5.2](#) above. This is simply another case where a strong mixing function is desired, mixing the input bits to produce a smaller number of output bits. The technique given in [Section 5.2.1](#) of using the parity of a number of bits is simply the result of successively Exclusive Or'ing them which is examined as a trivial mixing function immediately below. Use of stronger mixing functions to extract more of the

randomness in a stream of skewed bits is examined in [Section 6.1.2](#).

[6.1.1](#) A Trivial Mixing Function

A trivial example for single bit inputs described only for expository purposes is the Exclusive Or function, which is equivalent to addition without carry, as show in the table below. This is a degenerate case in which the one output bit always changes for a change in either input bit. But, despite its simplicity, it provides a useful illustration.

+-----+-----+-----+			
input 1	input 2	output	
+-----+-----+-----+			
0	0	0	
0	1	1	
1	0	1	
1	1	0	
+-----+-----+-----+			

If inputs 1 and 2 are uncorrelated and combined in this fashion then the output will be an even better (less skewed) random bit than the inputs. If we assume an "eccentricity" e as defined in [Section 5.2](#) above, then the output eccentricity relates to the input eccentricity as follows:

$$e_{\text{output}} = 2 * e_{\text{input 1}} * e_{\text{input 2}}$$

Since e is never greater than $1/2$, the eccentricity is always improved except in the case where at least one input is a totally skewed constant. This is illustrated in the following table where the top and left side values are the two input eccentricities and the entries are the output eccentricity:

+-----+-----+-----+-----+-----+-----+-----+

e	0.00	0.10	0.20	0.30	0.40	0.50
0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.10	0.00	0.02	0.04	0.06	0.08	0.10
0.20	0.00	0.04	0.08	0.12	0.16	0.20
0.30	0.00	0.06	0.12	0.18	0.24	0.30
0.40	0.00	0.08	0.16	0.24	0.32	0.40
0.50	0.00	0.10	0.20	0.30	0.40	0.50

However, keep in mind that the above calculations assume that the inputs are not correlated. If the inputs were, say, the parity of the number of minutes from midnight on two clocks accurate to a few seconds, then each might appear random if sampled at random intervals much longer than a minute. Yet if they were both sampled and combined with xor, the result would be zero most of the time.

[6.1.2](#) Stronger Mixing Functions

The US Government Advanced Encryption Standard [[AES](#)] is an example of a strong mixing function for multiple bit quantities. It takes up to 384 bits of input (128 bits of "data" and 256 bits of "key") and produces 128 bits of output each of which is dependent on a complex

D. Eastlake, J. Schiller, S. Crocker

[Page 23]

INTERNET DRAFT

Randomness Requirements for Security

October 2004

non-linear function of all input bits. Other encryption functions with this characteristic, such as [[DES](#)], can also be used by considering them to mix all of their key and data input bits.

Another good family of mixing functions are the "message digest" or hashing functions such as The US Government Secure Hash Standards [SHA*] and the MD4, MD5 [[MD4](#), [MD5](#)] series. These functions all take a practically unlimited amount of input and produce a relatively short fixed length output mixing all the input bits. The MD* series produce 128 bits of output, SHA-1 produces 160 bits, and other SHA functions produce up to 512 bits.

Although the message digest functions are designed for variable amounts of input, AES and other encryption functions can also be used to combine any number of inputs. If 128 bits of output is adequate, the inputs can be packed into a 128-bit data quantity and successive AES keys, padding with zeros if needed, which are then used to

successively encrypt using AES in Electronic Codebook Mode. Or the input could be packed into one 128-bit key and multiple data blocks and a CBC-MAC calculated [[MODES](#)].

If more than 128 bits of output are needed and you want to employ AES, use more complex mixing. But keep in mind that it is absolutely impossible to get more bits of "randomness" out than are put in. For example, if inputs are packed into three quantities, A, B, and C, use AES to encrypt A with B as a key and then with C as a key to produce the 1st part of the output, then encrypt B with C and then A for more output and, if necessary, encrypt C with A and then B for yet more output. Still more output can be produced by reversing the order of the keys given above to stretch things. The same can be done with the hash functions by hashing various subsets of the input data or different copies of the input data with different prefixes and/or suffixes to produce multiple outputs.

An example of using a strong mixing function would be to reconsider the case of a string of 308 bits each of which is biased 99% towards zero. The parity technique given in [Section 5.2.1](#) above reduced this to one bit with only a 1/1000 deviance from being equally likely a zero or one. But, applying the equation for information given in [Section 2](#), this 308 bit skewed sequence has over 5 bits of information in it. Thus hashing it with SHA-1 and taking the bottom 5 bits of the result would yield 5 unbiased random bits as opposed to the single bit given by calculating the parity of the string. Alternatively, for some applications, you could use the entire hash output to retain almost all of the entropy.

[6.1.3](#) Using S-Boxes for Mixing

Many modern block encryption functions, including DES and AES, incorporate modules known as S-Boxes (substitution boxes). These produce a smaller number of outputs from a larger number of inputs through a complex non-linear mixing function which would have the effect of concentrating limited entropy in the inputs into the output.

S-Boxes sometimes incorporate bent Boolean functions (functions of an even number of bits producing one output bit with maximum non-linearity). Looking at the output for all input pairs differing in any particular bit position, exactly half the outputs are different. An S-Box in which each output bit is produced by a bent function such that any linear combination of these functions is also a bent function is called a "perfect S-Box".

S-boxes and various repeated application or cascades of such boxes can be used for mixing. [SBOX*]

[6.1.4](#) Diffie-Hellman as a Mixing Function

Diffie-Hellman exponential key exchange is a technique that yields a shared secret between two parties that can be made computationally infeasible for a third party to determine even if they can observe all the messages between the two communicating parties. This shared secret is a mixture of initial quantities generated by each of the parties [[D-H](#)].

If these initial quantities are random and uncorrelated, then the shared secret combines that randomness, but, of course, can not produce more randomness than the size of the shared secret generated.

While this is true if the Diffie-Hellman computation is performed privately, if an adversary can observe either of the public keys and knows the modulus being used, they need only search through the space of the other secret key in order to be able to calculate the shared secret [[D-H](#)]. So, conservatively, it would be best to consider public Diffie-Hellman to produce a quantity whose guessability corresponds to the worst of the two inputs.

[6.1.5](#) Using a Mixing Function to Stretch Random Bits

While it is not necessary for a mixing function to produce the same or fewer bits than its inputs, mixing bits cannot "stretch" the amount of random unpredictability present in the inputs. Thus four

inputs of 32 bits each where there is 12 bits worth of

unpredictability (such as 4,096 equally probable values) in each input cannot produce more than 48 bits worth of unpredictable output. The output can be expanded to hundreds or thousands of bits by, for example, mixing with successive integers, but the clever adversary's search space is still 2^{48} possibilities. Furthermore, mixing to fewer bits than are input will tend to strengthen the randomness of the output.

The last table in [Section 6.1.1](#) shows that mixing a random bit with a constant bit with Exclusive Or will produce a random bit. While this is true, it does not provide a way to "stretch" one random bit into more than one. If, for example, a random bit is mixed with a 0 and then with a 1, this produces a two bit sequence but it will always be either 01 or 10. Since there are only two possible values, there is still only the one bit of original randomness.

[6.1.6](#) Other Factors in Choosing a Mixing Function

For local use, AES has the advantages that it has been widely tested for flaws, is reasonably efficient in software, and is widely documented and implemented with hardware and software implementations available all over the world including open source code. The SHA* family have had a little less study and tend to require more CPU cycles than AES but there is no reason to believe they are flawed. Both SHA* and MD5 were derived from the earlier MD4 algorithm. They all have source code available [SHA*, MD*]. Some signs of weakness have been found in MD4 and MD5. In particular, MD4 has only three rounds and there are several independent breaks of the first two or last two rounds. And some collisions have been found in MD5 output.

AES was selected by a robust, public, and international process. It and SHA* have been vouched for by the US National Security Agency (NSA) on the basis of criteria that mostly remain secret, as was DES. While this has been the cause of much speculation and doubt, investigation of DES over the years has indicated that NSA involvement in modifications to its design, which originated with IBM, was primarily to strengthen it. No concealed or special weakness has been found in DES. It is likely that the NSA modifications to MD4 to produce the SHA algorithms similarly strengthened these algorithms, possibly against threats not yet known in the public cryptographic community.

Where input lengths are unpredictable, hash algorithms are a little more convenient to use than block encryption algorithms since they are generally designed to accept variable length inputs. Block encryption algorithms generally require an additional padding algorithm to accommodate inputs that are not an even multiple of the

INTERNET DRAFT

Randomness Requirements for Security

October 2004

block size.

As of the time of this document, the authors know of no patent claims to the basic AES, DES, SHA*, MD4, and MD5 algorithms other than patents for which an irrevocable royalty free license has been granted to the world. There may, of course, be basic patents of which the authors are unaware or patents on implementations or uses or other relevant patents issued or to be issued.

[6.2](#) Non-Hardware Sources of Randomness

The best source of input for mixing would be a hardware randomness such as ring oscillators, disk drive timing, thermal noise, or radioactive decay. However, if that is not available, there are other possibilities. These include system clocks, system or input/output buffers, user/system/hardware/network serial numbers and/or addresses and timing, and user input. Unfortunately, each of these sources can produce very limited or predictable values under some circumstances.

Some of the sources listed above would be quite strong on multi-user systems where, in essence, each user of the system is a source of randomness. However, on a small single user or embedded system, especially at start up, it might be possible for an adversary to assemble a similar configuration. This could give the adversary inputs to the mixing process that were sufficiently correlated to those used originally as to make exhaustive search practical.

The use of multiple random inputs with a strong mixing function is recommended and can overcome weakness in any particular input. The timing and content of requested "random" user keystrokes can yield hundreds of random bits but conservative assumptions need to be made. For example, assuming at most a few bits of randomness if the inter-keystroke interval is unique in the sequence up to that point and a similar assumption if the key hit is unique but assuming that no bits of randomness are present in the initial key value or if the timing or key value duplicate previous values. The results of mixing these timings and characters typed could be further combined with clock values and other inputs.

This strategy may make practical portable code to produce good random numbers for security even if some of the inputs are very weak on some

of the target systems. However, it may still fail against a high grade attack on small, single user or embedded systems, especially if the adversary has ever been able to observe the generation process in the past. A hardware based random source is still preferable.

[6.3](#) Cryptographically Strong Sequences

In cases where a series of random quantities must be generated, an adversary may learn some values in the sequence. In general, they should not be able to predict other values from the ones that they know.

The correct technique is to start with a strong random seed, take cryptographically strong steps from that seed [[FERGUSON](#), [SCHNEIER](#)], and do not reveal the complete state of the generator in the sequence elements. If each value in the sequence can be calculated in a fixed way from the previous value, then when any value is compromised, all future values can be determined. This would be the case, for example, if each value were a constant function of the previously used values, even if the function were a very strong, non-invertible message digest function.

(It should be noted that if your technique for generating a sequence of key values is fast enough, it can trivially be used as the basis for a confidentiality system. If two parties use the same sequence generating technique and start with the same seed material, they will generate identical sequences. These could, for example, be xor'ed at one end with data being send, encrypting it, and xor'ed with this data as received, decrypting it due to the reversible properties of the xor operation. This is commonly referred to as a simple stream cipher.)


[6.3.1](#) OFB and CTR Sequences

One way to achieve a strong sequence is to have the values be produced by taking a seed value and hashing the quantities produced by concatenating the seed with successive integers or the like and then mask the values obtained so as to limit the amount of generator

state available to the adversary.

It may also be possible to use an "encryption" algorithm with a random key and seed value to encrypt successive integers as in counter (CTR) mode encryption. Alternatively, you can feedback all of the output encrypted value into the value to be encrypted for the next iteration. This is a particular example of output feedback mode (OFB). [[MODES](#)]

An example is shown below where shifting and masking are used to combine part of the output feedback with part of the old input. This type of partial feedback should be avoided for reasons described below.



Note that if a shift of one is used, this is the same as the shift register technique described in [Section 3](#) above but with the all important difference that the feedback is determined by a complex non-linear function of all bits rather than a simple linear or polynomial combination of output from a few bit position taps.

It has been shown by Donald W. Davies that this sort of shifted partial output feedback significantly weakens an algorithm compared with feeding all of the output bits back as input. In particular, for DES, repeated encrypting a full 64 bit quantity will give an expected repeat in about 2^{63} iterations. Feeding back anything less than 64 (and more than 0) bits will give an expected repeat in between 2^{31}

and 2^{32} iterations!

To predict values of a sequence from others when the sequence was generated by these techniques is equivalent to breaking the cryptosystem or inverting the "non-invertible" hashing involved with only partial information available. The less information revealed each iteration, the harder it will be for an adversary to predict the sequence. Thus it is best to use only one bit from each value. It has been shown that in some cases this makes it impossible to break a system even when the cryptographic system is invertible and can be broken if all of each generated value was revealed.

[6.3.2](#) The Blum Blum Shub Sequence Generator

Currently the generator which has the strongest public proof of strength is called the Blum Blum Shub generator after its inventors [BBS]. It is also very simple and is based on quadratic residues. Its only disadvantage is that it is computationally intensive compared with the traditional techniques give in 6.3.1 above. This is not a major draw back if it is used for moderately infrequent purposes, such as generating session keys.

D. Eastlake, J. Schiller, S. Crocker

[Page 29]

INTERNET DRAFT

Randomness Requirements for Security

October 2004

Simply choose two large prime numbers, say p and q , which both have the property that you get a remainder of 3 if you divide them by 4. Let $n = p * q$. Then you choose a random number x relatively prime to n . The initial seed for the generator and the method for calculating subsequent values are then

$$s_0 = (x^2) \pmod{n}$$

$$s_{i+1} = (s_i^2) \pmod{n}$$

You must be careful to use only a few bits from the bottom of each s . It is always safe to use only the lowest order bit. If you use no more than the

$$\log_2 (\log_2 (s_i))$$

low order bits, then predicting any additional bits from a sequence generated in this manner is provable as hard as factoring n . As long as the initial x is secret, you can even make n public if you want.


An interesting characteristic of this generator is that you can directly calculate any of the s values. In particular

$$s_i = (s_0^{((2^i \text{Mod } ((p-1) * (q-1))))}) \text{Mod } n$$

This means that in applications where many keys are generated in this fashion, it is not necessary to save them all. Each key can be effectively indexed and recovered from that small index and the initial s and n .

6.3.3 Entropy Pool Techniques

Many modern pseudo-random number sources utilize the technique of maintaining a "pool" of bits and providing operations for strongly mixing input with some randomness into the pool and extracting pseudo random bits from the pool. This is illustrated in the figure below.



Bits to be feed into the pool can be any of the various hardware, environmental, or user input sources discussed above. It is also common to save the state of the pool on system shut down and restore

it on re-starting, if stable storage is available.

Care must be taken that enough entropy has been added to the pool to support particular output uses desired. See [Section 7.5](#) for more details on an example implementation and [RSA BULL1] for similar suggestions.

[7.](#) Key Generation Examples and Standards

Several public standards and widely deployed examples are now in

place for the generation of keys without special hardware. Three standards are described below. The two older standards use DES, with its 64-bit block and key size limit, but any equally strong or stronger mixing function could be substituted [[DES](#)]. The third is a more modern and stronger standard based on SHA-1 [[SHA*](#)]. Lastly the widely deployed modern UNIX random number generators are described.

[7.1](#) US DoD Recommendations for Password Generation

The United States Department of Defense has specific recommendations for password generation [[DoD](#)]. They suggest using the US Data Encryption Standard [[DES](#)] in Output Feedback Mode [[MODES](#)] as follows:

- use an initialization vector determined from
 - the system clock,
 - system ID,
 - user ID, and
 - date and time;
- use a key determined from
 - system interrupt registers,
 - system status registers, and
 - system counters; and,
- as plain text, use an external randomly generated 64 bit quantity such as 8 characters typed in by a system administrator.

The password can then be calculated from the 64 bit "cipher text" generated by DES in 64-bit Output Feedback Mode. As many bits as are needed can be taken from these 64 bits and expanded into a pronounceable word, phrase, or other format if a human being needs to remember the password.

[7.2](#) X9.17 Key Generation

The American National Standards Institute has specified a method for generating a sequence of keys as follows [[X9.17](#)]:

s is the initial 64 bit seed
0

g is the sequence of generated 64 bit key quantities
n

k is a random key reserved for generating this key sequence

t is the time at which a key is generated to as fine a resolution as is available (up to 64 bits).

$DES (K, Q)$ is the DES encryption of quantity Q with key K

$$g_n = DES (k, DES (k, t) .xor. s_n)$$

$$s_{n+1} = DES (k, DES (k, t) .xor. g_n)$$

If $g_{sub\ n}$ is to be used as a DES key, then every eighth bit should be adjusted for parity for that use but the entire 64 bit unmodified g should be used in calculating the next s .

[7.3](#) DSS Pseudo-Random Number Generation

Appendix 3 of the NIST Digital Signature Standard [[DSS](#)] provides a method of producing a sequence of pseudo-random 160 bit quantities for use as private keys or the like. This has been modified by Change Notice 1 [DSS CN1] to produce the following algorithm for generating general purpose pseudorandom numbers:

$t = 0x\ 67452301\ EFCDAB89\ 98BADCFE\ 10325476\ C3D2E1F0$

$XKEY_0 = \text{initial seed}$

For $j = 0$ to ...

$$XVAL_j = (XKEY_j + \text{optional user input}) \text{ (Mod } 2^{512} \text{)}$$

$$X_j = G(t, XVAL_j)$$

$$XKEY_{j+1} = (1 + XKEY_j + X_j) \text{ (Mod } 2^{512} \text{)}$$

The quantities X thus produced are the pseudo-random sequence of 160 bit values. Two functions can be used for "G" above. Each produces a 160-bit value and takes two arguments, the first argument a 160-bit

value and the second a 512 bit value.

The first is based on SHA-1 and works by setting the 5 linking

variables, denoted H with subscripts in the SHA-1 specification, to the first argument divided into fifths. Then steps (a) through (e) of [section 7](#) of the NIST SHA-1 specification are run over the second argument as if it were a 512-bit data block. The values of the linking variable after those steps are then concatenated to produce the output of G. [SHA*]

As an alternative second method, NIST also defined an alternate G function based on multiple applications of the DES encryption function [[DSS](#)].

[7.4](#) X9.82 Pseudo-Random Number Generation

The National Institute for Standards and Technology (NIST) and the American National Standards Institutes (ANSI) X9F1 committee are in the final stages of creating a standard for random number generation covering both true randomness generators and pseudo-random number generators. It includes a number of pseudo-random number generators for use with AES and other block ciphers. It also includes random number generators based on hash functions and the arithmetic of elliptic curves [[X9.82](#)].

[7.5](#) The /dev/random Device

Several versions of the UNIX operating system provides a kernel-resident random number generator. In some cases, these generators makes use of events captured by the Kernel during normal system operation.

For example, on some versions of Linux, the generator consists of a random pool of 512 bytes represented as 128 words of 4-bytes each. When an event occurs, such as a disk drive interrupt, the time of the event is xor'ed into the pool and the pool is stirred via a primitive polynomial of degree 128. The pool itself is treated as a ring buffer, with new data being XORed (after stirring with the

polynomial) across the entire pool.

Each call that adds entropy to the pool estimates the amount of likely true entropy the input contains. The pool itself contains an accumulator that estimates the total over all entropy of the pool.

Input events come from several sources as listed below.

Unfortunately, for server machines without human operators, the first and third are not available and entropy may be added very slowly in that case.

1. Keyboard interrupts. The time of the interrupt as well as the scan code are added to the pool. This in effect adds entropy from the human operator by measuring inter-keystroke arrival times.
2. Disk completion and other interrupts. A system being used by a person will likely have a hard to predict pattern of disk accesses. (But not all disk drivers support capturing this timing information with sufficient accuracy to be useful.)
3. Mouse motion. The timing as well as mouse position is added in.

When random bytes are required, the pool is hashed with SHA-1 [SHA*] to yield the returned bytes of randomness. If more bytes are required than the output of SHA-1 (20 bytes), then the hashed output is stirred back into the pool and a new hash performed to obtain the next 20 bytes. As bytes are removed from the pool, the estimate of entropy is similarly decremented.

To ensure a reasonable random pool upon system startup, the standard startup scripts (and shutdown scripts) save the pool to a disk file at shutdown and read this file at system startup.

There are two user exported interfaces. /dev/random returns bytes from the pool, but blocks when the estimated entropy drops to zero. As entropy is added to the pool from events, more data becomes available via /dev/random. Random data obtained from such a /dev/random device is suitable for key generation for long term keys, if enough random bits are in the pool or are added in a reasonable amount of time.

/dev/urandom works like /dev/random, however it provides data even

when the entropy estimate for the random pool drops to zero. This may be adequate for session keys or for other key generation tasks where blocking while waiting for more random bits is not acceptable. The risk of continuing to take data even when the pool's entropy estimate is small is that past output may be computable from current output provided an attacker can reverse SHA-1. Given that SHA-1 is designed to be non-invertible, this is a reasonable risk.

To obtain random numbers under Linux, Solaris, or other UNIX systems equipped with code as described above, all an application needs to do is open either `/dev/random` or `/dev/urandom` and read the desired number of bytes.

(The Linux Random device was written by Theodore Ts'o. It was based loosely on the random number generator in PGP 2.X and PGP 3.0 (aka PGP 5.0).)

[7.6](#) Windows CryptGenRandom

Microsoft's recommendation to users of the widely deployed Windows operating system is generally to use the CryptGenRandom pseudo-random number generation call with the CryptAPI cryptographic service provider. This takes a handle to a cryptographic service provider library, a pointer to a buffer by which the caller can provide entropy and into which the generated pseudo-randomness is returned, and an indication of how many octets of randomness are desired.

The Windows CryptAPI cryptographic service provider stores a seed state variable with every user. When CryptGenRandom is called, this is combined with any randomness provided in the call and various system and user data such as the process ID, thread ID, system clock, system time, system counter, memory status, free disk clusters, and hashed user environment block. This data is all feed to SHA-1 and the output used to seed an RC4 key stream. That key stream is used to produce the pseudo-random data requested and to update the user's seed state variable.

Users of Windows ".NET" will probably find it easier to use the `RNGCryptoServiceProvider.GetBytes` method interface.

For further information, see [[WSC](#)].

[8](#). Examples of Randomness Required

Below are two examples showing rough calculations of needed randomness for security. The first is for moderate security passwords while the second assumes a need for a very high security cryptographic key.

In addition [[ORMAN](#)] and [RSA BULL13] provide information on the public key lengths that should be used for exchanging symmetric keys.

[8.1](#) Password Generation

Assume that user passwords change once a year and it is desired that

the probability that an adversary could guess the password for a particular account be less than one in a thousand. Further assume that sending a password to the system is the only way to try a password. Then the crucial question is how often an adversary can try possibilities. Assume that delays have been introduced into a system so that, at most, an adversary can make one password try every six seconds. That's 600 per hour or about 15,000 per day or about 5,000,000 tries in a year. Assuming any sort of monitoring, it is unlikely someone could actually try continuously for a year. In fact, even if log files are only checked monthly, 500,000 tries is more plausible before the attack is noticed and steps taken to change passwords and make it harder to try more passwords.

To have a one in a thousand chance of guessing the password in 500,000 tries implies a universe of at least 500,000,000 passwords or about 2^{29} . Thus 29 bits of randomness are needed. This can probably be achieved using the US DoD recommended inputs for password generation as it has 8 inputs which probably average over 5 bits of randomness each (see [section 7.1](#)). Using a list of 1000 words, the password could be expressed as a three word phrase (1,000,000,000 possibilities) or, using case insensitive letters and digits, six would suffice ($(26+10)^6 = 2,176,782,336$ possibilities).

For a higher security password, the number of bits required goes up. To decrease the probability by 1,000 requires increasing the universe of passwords by the same factor which adds about 10 bits. Thus to have only a one in a million chance of a password being guessed under the above scenario would require 39 bits of randomness and a password that was a four word phrase from a 1000 word list or eight letters/digits. To go to a one in 10^9 chance, 49 bits of randomness are needed implying a five word phrase or ten letter/digit password.

In a real system, of course, there are also other factors. For example, the larger and harder to remember passwords are, the more likely users are to write them down resulting in an additional risk

of compromise.

[8.2](#) A Very High Security Cryptographic Key

Assume that a very high security key is needed for symmetric

encryption / decryption between two parties. Assume an adversary can observe communications and knows the algorithm being used. Within the field of random possibilities, the adversary can try key values in hopes of finding the one in use. Assume further that brute force trial of keys is the best the adversary can do.

8.2.1 Effort per Key Trial

How much effort will it take to try each key? For very high security applications it is best to assume a low value of effort. Even if it would clearly take tens of thousands of computer cycles or more to try a single key, there may be some pattern that enables huge blocks of key values to be tested with much less effort per key. Thus it is probably best to assume no more than a couple hundred cycles per key. (There is no clear lower bound on this as computers operate in parallel on a number of bits and a poor encryption algorithm could allow many keys or even groups of keys to be tested in parallel. However, we need to assume some value and can hope that a reasonably strong algorithm has been chosen for our hypothetical high security task.)

If the adversary can command a highly parallel processor or a large network of work stations, 10^{11} cycles per second is probably a minimum assumption for availability today. Looking forward a few years, there should be at least an order of magnitude improvement. Thus assuming 10^{10} keys could be checked per second or 3.6×10^{12} per hour or 6×10^{14} per week or 2.4×10^{15} per month is reasonable. This implies a need for a minimum of 63 bits of randomness in keys to be sure they cannot be found in a month. Even then it is possible that, a few years from now, a highly determined and resourceful adversary could break the key in 2 weeks (on average they need try only half the keys).

These questions are considered in detail in "Minimal Key Lengths for Symmetric Ciphers to Provide Adequate Commercial Security: A Report by an Ad Hoc Group of Cryptographers and Computer Scientists" [[KeyStudy](#)] which was sponsored by the Business Software Alliance. It concluded that a reasonable key length in 1995 for very high security is in the range of 75 to 90 bits and, since the cost of cryptography does not vary much with the key size, recommends 90 bits. To update these recommendations, just add 2/3 of a bit per year for Moore's law

[[MOORE](#)]. Thus, in the year 2004, this translates to a determination that a reasonable key length is in the 81 to 96 bit range. In fact, today, it is increasingly common to use keys longer than 96 bits, such as 128-bit (or longer) keys with AES and keys with effective lengths of 112-bits using triple-DES.

[8.2.2](#) Meet in the Middle Attacks

If chosen or known plain text and the resulting encrypted text are available, a "meet in the middle" attack is possible if the structure of the encryption algorithm allows it. (In a known plain text attack, the adversary knows all or part of the messages being encrypted, possibly some standard header or trailer fields. In a chosen plain text attack, the adversary can force some chosen plain text to be encrypted, possibly by "leaking" an exciting text that would then be sent by the adversary over an encrypted channel.)

An oversimplified explanation of the meet in the middle attack is as follows: the adversary can half-encrypt the known or chosen plain text with all possible first half-keys, sort the output, then half-decrypt the encoded text with all the second half-keys. If a match is found, the full key can be assembled from the halves and used to decrypt other parts of the message or other messages. At its best, this type of attack can halve the exponent of the work required by the adversary while adding a very large but roughly constant factor of effort. Thus, if this attack can be mounted, a doubling of the amount of randomness in the very strong key to a minimum of 192 bits (96×2) is required for the year 2004 based on the [[KeyStudy](#)] analysis.

This amount of randomness is well beyond the limit of that in the inputs recommended by the US DoD for password generation and could require user typing timing, hardware random number generation, or other sources.

The meet in the middle attack assumes that the cryptographic algorithm can be decomposed in this way. Hopefully no modern algorithm has this weakness but there may be cases where we are not sure of that or even of what algorithm a key will be used with. Even if a basic algorithm is not subject to a meet in the middle attack, an attempt to produce a stronger algorithm by applying the basic algorithm twice (or two different algorithms sequentially) with different keys will gain less added security than would be expected. Such a composite algorithm would be subject to a meet in the middle attack.

Enormous resources may be required to mount a meet in the middle attack but they are probably within the range of the national

security services of a major nation. Essentially all nations spy on other nations traffic.

[8.2.3](#) Other Considerations

[KeyStudy] also considers the possibilities of special purpose code breaking hardware and having an adequate safety margin.

It should be noted that key length calculations such as those above are controversial and depend on various assumptions about the cryptographic algorithms in use. In some cases, a professional with a deep knowledge of code breaking techniques and of the strength of the algorithm in use could be satisfied with less than half of the 192 bit key size derived above.

For further examples of conservative design principles see [\[FERGUSON\]](#).

9. Conclusion

Generation of unguessable "random" secret quantities for security use is an essential but difficult task.

Hardware techniques to produce such randomness would be relatively simple. In particular, the volume and quality would not need to be high and existing computer hardware, such as audio input or disk drives, can be used.

Widely available computational techniques are available to process low quality random quantities from multiple sources or a larger quantity of such low quality input from one source and produce a smaller quantity of higher quality keying material. In the absence of hardware sources of randomness, a variety of user and software sources can frequently, with care, be used instead; however, most modern systems already have hardware, such as disk drives or audio input, that could be used to produce high quality randomness.

Once a sufficient quantity of high quality seed key material (a couple of hundred bits) is available, computational techniques are available to produce cryptographically strong sequences of computationally unpredictable quantities from this seed material.

[10](#). Security Considerations

The entirety of this document concerns techniques and recommendations for generating unguessable "random" quantities for use as passwords, cryptographic keys, initialization vectors, sequence numbers, and similar security uses.

[11](#). Copyright and Disclaimer

Copyright (C) The Internet Society 2004. This document is subject to the rights, licenses and restrictions contained in [BCP 78](#) and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

[12. Appendix A](#): Changes from [RFC 1750](#)

1. Additional acknowledgements have been added.
2. Insertion of [section 5.2.4](#) on de-skewing with S-boxes.
3. Addition of [section 5.4](#) on Ring Oscillator randomness sources.
4. AES and the members of the SHA series producing more than 160 bits have been added. Use of AES has been emphasized and the use of DES de-emphasized.
5. Addition of [section 6.3.3](#) on entropy pool techniques.
6. Addition of [section 7.3](#) on the pseudo-random number generation techniques given in FIPS 186-2 (with Change Notice 1), 7.4 on those given in X9.82, [section 7.5](#) on the random number generation techniques of the /dev/random device in Linux and other UNIX

systems, and [section 7.6](#) on random number generation techniques in the Windows operating system.

7. Addition of references to the "Minimal Key Lengths for Symmetric Ciphers to Provide Adequate Commercial Security" study published in January 1996 [[KeyStudy](#)].
8. Added caveats to using Diffie-Hellman as a mixing function.
9. Addition of references to the [[TURBID](#)] paper and system.
10. Addition of discussion of min-entropy and Renyi entropy and references to the [[LUBY](#)] book.
11. Minor wording changes and reference updates.

[14](#). Informative References

[AES] - "Specification of the Advanced Encryption Standard (AES)", United States of America, US National Institute of Standards and Technology, FIPS 197, November 2001.

[ASYMMETRIC] - "Secure Communications and Asymmetric Cryptosystems", edited by Gustavus J. Simmons, AAAS Selected Symposium 69, Westview Press, Inc.

[BBS] - "A Simple Unpredictable Pseudo-Random Number Generator", SIAM Journal on Computing, v. 15, n. 2, 1986, L. Blum, M. Blum, & M. Shub.

[BRILLINGER] - "Time Series: Data Analysis and Theory", Holden-Day, 1981, David Brillinger.

[CRC] - "C.R.C. Standard Mathematical Tables", Chemical Rubber Publishing Company.

[DAVIS] - "Cryptographic Randomness from Air Turbulence in Disk Drives", Advances in Cryptology - Crypto '94, Springer-Verlag Lecture Notes in Computer Science #839, 1984, Don Davis, Ross Ihaka, and Philip Fenstermacher.

[DES] - "Data Encryption Standard", US National Institute of Standards and Technology, FIPS 46-3, October 1999.

- "Data Encryption Algorithm", American National Standards Institute, ANSI X3.92-1981.

(See also FIPS 112, Password Usage, which includes FORTRAN code for performing DES.)

[D-H] - [RFC 2631](#), "Diffie-Hellman Key Agreement Method", Eric Rescorla, June 1999.

[DNSSEC] - [RFC 2535](#), "Domain Name System Security Extensions", D. Eastlake, March 1999.

[DoD] - "Password Management Guideline", United States of America, Department of Defense, Computer Security Center, CSC-STD-002-85.
(See also FIPS 112, Password Usage, which incorporates CSC-STD-002-85 as one of its appendices.)

[DSS] - "Digital Signature Standard (DSS)", US National Institute of Standards and Technology, FIPS 186-2, January 2000.

[DSS CN1] - "Digital Signature Standard Change Notice 1", US National Institute of Standards and Technology, FIPS 186-2 Change Notice 1, 5 October 2001.

[FERGUSON] - "Practical Cryptography", Niels Ferguson and Bruce

[GIFFORD] - "Natural Random Number", MIT/LCS/TM-371, David K. Gifford, September 1988.

[IEEE 802.11i] - "Amendment to Standard for Telecommunications and Information Exchange Between Systems - LAN/MAN Specific Requirements - Part 11: Wireless Medium Access Control (MAC) and physical layer (PHY) specifications: Medium Access Control (MAC) Security Enhancements", The Institute for Electrical and Electronics Engineers, January 2004.

[IPSEC] - [RFC 2401](#), "Security Architecture for the Internet Protocol", S. Kent, R. Atkinson, November 1998.

[Jakobsson] - M. Jakobsson, E. Shriver, B. K. Hillyer, and A. Juels, "A practical secure random bit generator", Proceedings of the Fifth ACM Conference on Computer and Communications Security, 1998. See also <http://citeseer.ist.psu.edu/article/jakobsson98practical.html>.

[KAUFMAN] - "Network Security: Private Communication in a Public World", Charlie Kaufman, Radia Perlman, and Mike Speciner, Prentis Hall PTR, ISBN 0-13-046019-2, 2nd Edition 2002.

[KeyStudy] - "Minimal Key Lengths for Symmetric Ciphers to Provide Adequate Commercial Security: A Report by an Ad Hoc Group of Cryptographers and Computer Scientists", M. Blaze, W. Diffie, R. Rivest, B. Schneier, T. Shimomura, E. Thompson, and M. Weiner, January 1996, <www.counterpane.com/keylength.html>.

[KNUTH] - "The Art of Computer Programming", Volume 2: Seminumerical Algorithms, Chapter 3: Random Numbers, Donald E. Knuth, Addison Wesley Publishing Company, 3rd Edition November 1997.

[KRAWCZYK] - "How to Predict Congruential Generators", H. Krawczyk, Journal of Algorithms, V. 13, N. 4, December 1992.

[LUBY] - "Pseudorandomness and Cryptographic Applications", Michael Luby, Princeton University Press, ISBN 0691025460, 8 January 1996.

[MAIL PEM 1] - [RFC 1421](#), "Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", J. Linn, 02/10/1993.

[MAIL PEM 2] - [RFC 1422](#), "Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management", S. Kent, 02/10/1993.

[MAIL PEM 3] - [RFC 1423](#), "Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers", D. Balenson, 02/10/1993.

[MAIL PEM 4] - [RFC 1424](#), "Privacy Enhancement for Internet

Electronic Mail: Part IV: Key Certification and Related Services", B. Kaliski, 02/10/1993.

[MAIL PGP]

- [RFC 2440](#), "OpenPGP Message Format", J. Callas, L. Donnerhackle, H. Finney, R. Thayer", November 1998.
- [RFC 3156](#), "MIME Security with OpenPGP" M. Elkins, D. Del Torto, R. Levien, T. Roessler, August 2001.

[MAIL S/MIME] - RFCs 2632 through 2634:

- [RFC 2632](#), "S/MIME Version 3 Certificate Handling", B. Ramsdell, Ed., June 1999.
- [RFC 2633](#), "S/MIME Version 3 Message Specification", B. Ramsdell, Ed., June 1999.
- [RFC 2634](#), "Enhanced Security Services for S/MIME" P. Hoffman, Ed., June 1999.

[MD4] - "The MD4 Message-Digest Algorithm", [RFC1320](#), April 1992, R. Rivest

[MD5] - "The MD5 Message-Digest Algorithm", [RFC1321](#), April 1992, R. Rivest

[MODES] - "DES Modes of Operation", US National Institute of Standards and Technology, FIPS 81, December 1980.

- "Data Encryption Algorithm - Modes of Operation", American National Standards Institute, ANSI X3.106-1983.

[MOORE] - Moore's Law: the exponential increase in the logic density of silicon circuits. Originally formulated by Gordon Moore in 1964 as a doubling every year starting in 1962, in the late 1970s the rate fell to a doubling every 18 months and has remained there through the date of this document. See "The New Hacker's Dictionary", Third Edition, MIT Press, ISBN 0-262-18178-9, Eric S. Raymond, 1996.

[ORMAN] - "Determining Strengths For Public Keys Used For Exchanging Symmetric Keys", [RFC 3766](#), Hilarie Orman, Paul Hoffman, April 2004.

[RFC 1750] - "Randomness Requirements for Security", D. Eastlake, S. Crocker, J. Schiller, December 1994.

[RSA BULL1] - "Suggestions for Random Number Generation in Software", RSA Laboratories Bulletin #1, January 1996.

[RSA BULL13] - "A Cost-Based Security Analysis of Symmetric and Asymmetric Key Lengths", RSA Laboratories Bulletin #13, Robert

Silverman, April 2000 (revised November 2001).

[SB0X1] - "Practical s-box design", S. Mister, C. Adams, Selected Areas in Cryptography, 1996.

D. Eastlake, J. Schiller, S. Crocker

[Page 46]

INTERNET DRAFT

Randomness Requirements for Security

October 2004

[SB0X2] - "Perfect Non-linear S-boxes", K. Nyberg, Advances in Cryptography - Eurocrypt '91 Proceedings, Springer-Verland, 1991.

[SCHNEIER] - "Applied Cryptography: Protocols, Algorithms, and Source Code in C", Bruce Schneier, 2nd Edition, John Wiley & Sons, 1996.

[SHANNON] - "The Mathematical Theory of Communication", University of Illinois Press, 1963, Claude E. Shannon. (originally from: Bell System Technical Journal, July and October 1948)

[SHIFT1] - "Shift Register Sequences", Solomon W. Golomb, Aegean Park Press, Revised Edition 1982.

[SHIFT2] - "Cryptanalysis of Shift-Register Generated Stream Cypher Systems", Wayne G. Barker, Aegean Park Press, 1984.

[SHA] - "Secure Hash Standard", US National Institute of Science and Technology, FIPS 180-2, 1 August 2002.

[SHA RFC] - [RFC 3174](#), "US Secure Hash Algorithm 1 (SHA1)", D. Eastlake, P. Jones, September 2001.

[SSH] - [draft-ietf-secsh](#)*, work in progress.

[STERN] - "Secret Linear Congruential Generators are not Cryptographically Secure", J. Stern, Proceedings of IEEE STOC, 1987.

[TLS] - [RFC 2246](#), "The TLS Protocol Version 1.0", T. Dierks, C. Allen, January 1999.

[TURBID] - "High Entropy Symbol Generator", John S. Denker, <<http://www.av8n.com/turbid/paper/turbid.htm>>, 2003.

[USENET] - [RFC 977](#), "Network News Transfer Protocol", B. Kantor, P. Lapsley, February 1986.

- [RFC 2980](#), "Common NNTP Extensions", S. Barber, October 2000.

[VON NEUMANN] - "Various techniques used in connection with random digits", von Neumann's Collected Works, Vol. 5, Pergamon Press, 1963, J. von Neumann.

[WSC] - "Writing Secure Code, Second Edition", Michael Howard, David. C. LeBlanc, Microsoft Press, ISBN 0735617228, December 2002.

[X9.17] - "American National Standard for Financial Institution Key Management (Wholesale)", American Bankers Association, 1985.

[X9.82] - "Random Number Generation", American National Standards Institute, ANSI X9F1, work in progress.

D. Eastlake, J. Schiller, S. Crocker

[Page 47]

INTERNET DRAFT Randomness Requirements for Security October 2004

Author's Addresses

Donald E. Eastlake 3rd
Motorola Laboratories
155 Beaver Street
Milford, MA 01757 USA

Telephone: +1 508-786-7554 (w)
 +1 508-634-2066 (h)
EMail: Donald.Eastlake@motorola.com

Jeffrey I. Schiller
MIT, Room E40-311
77 Massachusetts Avenue
Cambridge, MA 02139-4307 USA

Telephone: +1 617-253-0161
E-mail: jis@mit.edu

Steve Crocker

EMail: steve@stevecrocker.com

File Name and Expiration

This is file [draft-eastlake-randomness2-09.txt](#).

It expires April 2005.