

Minimal Key Lengths for Symmetric Ciphers to Provide Adequate Commercial Security

A Report by an Ad Hoc Group of
Cryptographers and Computer Scientists

Matt Blaze ¹
Whitfield Diffie ²
Ronald L. Rivest ³
Bruce Schneier ⁴
Tsutomu Shimomura ⁵
Eric Thompson ⁶
Michael Wiener ⁷

January 1996

ABSTRACT

Encryption plays an essential role in protecting the privacy of electronic information against threats from a variety of potential attackers. In so doing, modern cryptography employs a combination of *conventional* or *symmetric* cryptographic systems for encrypting data and *public key* or *asymmetric* systems for managing the *keys* used by the symmetric systems. Assessing the strength required of the symmetric cryptographic systems is therefore an essential step in employing cryptography for computer and communication security.

Technology readily available today (late 1995) makes *brute-force* attacks against cryptographic systems considered adequate for the past several years both fast and cheap. General purpose computers can be used, but a much more efficient approach is to employ commercially available *Field Programmable Gate Array (FPGA)* technology. For attackers prepared to make a higher initial investment, custom-made, special-purpose chips make such calculations much faster and significantly lower the amortized cost per solution.

As a result, cryptosystems with 40-bit keys offer virtually no protection at this point against brute-force attacks. Even the U.S. Data Encryption Standard with 56-bit keys is increasingly inadequate. As cryptosystems often succumb to ‘smarter’ attacks than brute-force key search, it is also important to remember that the keylengths discussed here are the minimum needed for security against the computational threats considered.

Fortunately, the cost of very strong encryption is not significantly greater than that of weak encryption. Therefore, to provide adequate protection against the most serious threats — well-funded commercial enterprises or government intelligence agencies — keys used to protect data today should be at least 75 bits long. To protect information adequately for the next 20 years in the face of expected advances in computing power, keys in newly-deployed systems should be at least 90 bits long.

¹AT&T Research, mab@research.att.com

²Sun Microsystems, diffie@eng.sun.com

³MIT Laboratory for Computer Science, rivest@lcs.mit.edu

⁴Counterpane Systems, schneier@counterpane.com

⁵San Diego Supercomputer Center, tsutomu@sdsc.edu

⁶Access Data, Inc., eric@accessdata.com

⁷Bell Northern Research, wieners@bnr.ca

1 Encryption Plays an Essential Role in Protecting the Privacy of Electronic Information

1.1 There is a need for information security.

As we write this paper in late 1995, the development of electronic commerce and the Global Information Infrastructure is at a critical juncture. The dirt paths of the middle ages only became highways of business and culture after the security of travelers and the merchandise they carried could be assured. So too the information superhighway will be an ill-traveled road unless information, the goods of the Information Age, can be moved, stored, bought, and sold securely. Neither corporations nor individuals will entrust their private business or personal data to computer networks unless they can assure their information's security.

Today, most forms of information can be stored and processed electronically. This means a wide variety of information, with varying economic values and privacy aspects and with a wide variation in the time over which the information needs to be protected, will be found on computer networks. Consider the spectrum:

- Electronic Funds Transfers of millions or even billions of dollars, whose short term security is essential but whose exposure is brief;
- A company's strategic corporate plans, whose confidentiality must be preserved for a small number of years;
- A proprietary product (Coke formula, new drug design) that needs to be protected over its useful life, often decades; and
- Information private to an individual (medical condition, employment evaluation) that may need protection for the lifetime of the individual.

1.2 Encryption can provide strong confidentiality protection.

Encryption is accomplished by scrambling data using mathematical procedures that make it extremely difficult and time consuming for anyone other than authorized recipients — those with the correct decryption *keys* — to recover the *plain text*. Proper encryption guarantees that the information will be safe even if it falls into hostile hands.

Encryption — and decryption — can be performed by either computer software or hardware. Common approaches include writing the algorithm on a disk for execution by a computer central processor; placing it in ROM or PROM for execution by a microprocessor; and isolating storage and execution in a computer accessory device (smart card or PCMCIA card).

The degree of protection obtained depends on several factors. These include: the quality of the cryptosystem; the way it is implemented in software or hardware (especially its reliability and the manner in which the keys are chosen); and the total number of possible keys that can be used to encrypt the information. A cryptographic algorithm is considered strong if:

1. There is no shortcut that allows the opponent to recover the plain text without using brute force to test keys until the correct one is found; and
2. The number of possible keys is sufficiently large to make such an attack infeasible.

The principle here is similar to that of a combination lock on a safe. If the lock is well designed so that a burglar cannot hear or feel its inner workings, a person who does not know the combination can open it only by dialing one set of numbers after another until it yields.

The sizes of encryption keys are measured in bits and the difficulty of trying all possible keys grows exponentially with the number of bits used. Adding one bit to the key doubles the number of possible keys; adding ten increases it by a factor of more than a thousand.

There is no definitive way to look at a cipher and determine whether a shortcut exists. Nonetheless, several encryption algorithms — most notably the U.S Data Encryption Standard (DES) — have been extensively studied in the public literature and are widely believed to be of very high quality. An essential element in cryptographic algorithm design is thus the length of the key, whose size places an upper bound on the system's strength.

Throughout this paper, we will assume that there are no shortcuts and treat the length of the key as representative of the cryptosystem's *workfactor* — the minimum amount of effort required to break the system. It is important to bear in mind, however, that cryptographers regard this as a rash assumption and many would recommend keys two or more times as long as needed to resist brute-force attacks. Prudent cryptographic designs not only employ longer keys than might appear to be needed, but devote more computation to encrypting and decrypting. A good example of this is the popular approach of using *triple-DES*: encrypting the output of DES twice more, using a total of three distinct keys.

Encryption systems fall into two broad classes. Conventional or symmetric cryptosystems — those in which an entity with the ability to encrypt also has the ability to decrypt and vice versa — are the systems under consideration in this paper. The more recent public key or asymmetric cryptosystems have the property that the ability to encrypt does not imply the ability to decrypt. In contemporary cryptography, public-key systems are indispensable for managing the keys of conventional cryptosystems. All known public key cryptosystems, however, are subject to shortcut attacks and must therefore use keys ten or more times the lengths of those discussed here to achieve the an equivalent level of security.

Although computers permit electronic information to be encrypted using very large keys, advances in computing power keep pushing up the size of keys that can be considered large and thus keep making it easier for individuals and organizations to attack encrypted information without the expenditure of unreasonable resources.

1.3 There are threats from a variety of potential attackers.

Threats to confidentiality of information come from a number of directions and their forms depend on the resources of the attackers. ‘Hackers,’ who might be anything from high school students to commercial programmers, may have access to mainframe computers or networks of workstations. The same people can readily buy inexpensive, off-the-shelf, boards, containing *Field Programmable Gate Array (FPGA)* chips that function as ‘programmable hardware’ and vastly increase the effectiveness of a cryptanalytic effort. A startup company or even a well-heeled individual could afford large numbers of these chips. A major corporation or organized crime operation with ‘serious money’ to spend could acquire custom computer chips specially designed for decryption. An intelligence agency, engaged in espionage for national economic advantage, could build a machine employing millions of such chips.

1.4 Current technology permits very strong encryption for effectively the same cost as weaker encryption.

It is a property of computer encryption that modest increases in computational cost can produce vast increases in security. Encrypting information very securely (e.g., with 128-bit keys) typically requires little more computing than encrypting it weakly (e.g., with 40-bit keys). In many applications, the cryptography itself accounts for only a small fraction of the computing costs, compared to such processes as voice or image compression required to prepare material for encryption.

One consequence of this uniformity of costs is that there is rarely any need to tailor the strength of cryptography to the sensitivity of the information being protected. Even if most of the information in a system has neither privacy implications nor monetary value, there is no practical or economic reason to design computer hardware or software to provide differing levels of encryption for different messages. It is simplest, most prudent, and thus fundamentally most economical, to employ a uniformly high level of encryption: the strongest encryption required for any information that might be stored or transmitted by a secure system.

2 Readily Available Technology Makes Brute-Force Decryption Attacks Faster and Cheaper

The kind of hardware used to mount a brute-force attack against an encryption algorithm depends on the scale of the cryptanalytic operation and the total funds available to the attacking enterprise. In the analysis that follows, we consider three general classes of technology that are likely to be employed by attackers with differing resources available to them. Not surprisingly, the cryptanalytic technologies that require larger up-front investments yield the lowest cost per recovered key, amortized over the life of the hardware.

It is the nature of brute-force attacks that they can be parallelized indefinitely. It is possible to use as many machines as are available, assigning each to work on a separate part of the problem. Thus regardless of the technology employed, the search time can be reduced by adding more equipment; twice as much hardware can be expected to find the right key in half the time. The total investment will have doubled, but if the hardware is kept constantly busy finding keys, the cost per key recovered is unchanged.

At the low end of the technology spectrum is the use of conventional personal computers or workstations programmed to test keys. Many people, by virtue of already owning or having access to the machines, are in a position use such resources at little or no cost. However, general purpose computers — laden with such ancillary equipment as video controllers, keyboards, interfaces, memory, and disk storage — make expensive search engines. They are therefore likely to be employed only by casual attackers who are unable or unwilling to invest in more specialized equipment.

A more efficient technological approach is to take advantage of commercially available Field Programmable Gate Arrays. FPGAs function as programmable hardware and allow faster implementations of such tasks as encryption and decryption than conventional processors. FPGAs are a commonly used tool for simple computations that need to be done very quickly, particularly simulating integrated circuits during development.

FPGA technology is fast and cheap. The cost of an AT&T ORCA chip that can test 30 million DES keys per second is \$200. This is 1,000 times faster than a PC at about one-tenth the cost! FPGAs are widely available and, mounted on cards, can be installed in standard PCs just like

sound cards, modems, or extra memory.

FPGA technology may be optimal when the same tool must be used for attacking a variety of different cryptosystems. Often, as with DES, a cryptosystem is sufficiently widely used to justify the construction of more specialized facilities. In these circumstances, the most cost-effective technology, but the one requiring the largest initial investment, is the use of *Application-Specific Integrated Circuits (ASICs)*. A \$10 chip can test 200 million keys per second. This is seven times faster than an FPGA chip at one-twentieth the cost.

Because ASICs require a far greater engineering investment than FPGAs and must be fabricated in quantity before they are economical, this approach is only available to serious, well-funded operations such as dedicated commercial (or criminal) enterprises and government intelligence agencies.

3 40-Bit Key Lengths Offer Virtually No Protection

Current U.S. Government policy generally limits exportable mass market software that incorporates encryption for confidentiality to using the RC2 or RC4 algorithms with 40-bit keys. A 40-bit key length means that there are 2^{40} possible keys. On average, half of these (2^{39}) must be tried to find the correct one. Export of other algorithms and key lengths must be approved on a case by case basis. For example, DES with a 56-bit key has been approved for certain applications such as financial transactions.

The recent successful brute-force attack by two French graduate students on Netscape's 40-bit RC4 algorithm demonstrates the dangers of such short keys. These students at the Ecole Polytechnique in Paris used 'idle time' on the school's computers, incurring no cost to themselves or their school. Even with these limited resources, they were able to recover the 40-bit key in a few days.

There is no need to have the resources of an institution of higher education at hand, however. Anyone with a modicum of computer expertise and a few hundred dollars would be able to attack 40-bit encryption much faster. An FPGA chip — costing approximately \$400 mounted on a card — would on average recover a 40-bit key in five hours. Assuming the FPGA lasts three years and is used continuously to find keys, the average cost per key is eight cents.

A more determined commercial predator, prepared to spend \$10,000 for a set-up with 25 ORCA chips, can find 40-bit keys in an average of 12 minutes, at the same average eight cent cost. Spending more money to buy more chips reduces the time accordingly: \$300,000 results in a solution in an average of 24 seconds; \$10,000,000 results in an average solution in 0.7 seconds.

As already noted, a corporation with substantial resources can design and commission custom chips that are much faster. By doing this, a company spending \$300,000 could find the right 40-bit key in an average of 0.18 seconds at 1/10th of a cent per solution; a larger company or government agency willing to spend \$10,000,000 could find the right key on average in 0.005 seconds (again at 1/10th of a cent per solution). (Note that the cost per solution remains constant because we have conservatively assumed constant costs for chip acquisition — in fact increasing the quantities purchased of a custom chip reduces the average chip cost as the initial design and set-up costs are spread over a greater number of chips.)

These results are summarized in Table I.

4 Even DES with 56-Bit Keys Is Increasingly Inadequate

4.1 DES is no panacea today.

The Data Encryption Standard (DES) was developed in the 1970s by IBM and NSA and adopted by the U.S. Government as a Federal Information Processing Standard for data encryption. It was intended to provide strong encryption for the government's sensitive but unclassified information. It was recognized by many, even at the time DES was adopted, that technological developments would make DES's 56-bit key exceedingly vulnerable to attack before the end of the century.

Today, DES may be the most widely employed encryption algorithm and continues to be a commonly cited benchmark. Yet DES-like encryption strength is no panacea. Calculations show that DES is inadequate against a corporate or government attacker committing serious resources. The bottom line is that DES is cheaper and easier to break than many believe.

As explained above, 40-bit encryption provides inadequate protection against even the most casual of intruders, content to scavenge time on idle machines or to spend a few hundred dollars. Against such opponents, using DES with a 56-bit key will provide a substantial measure of security. At present, it would take a year and a half for someone using \$10,000 worth of FPGA technology to search out a DES key. In ten years time an investment of this size would allow one to find a DES key in less than a week.

The real threat to commercial transactions and to privacy on the Internet is from individuals and organizations willing to invest substantial time and money. As more and more business and personal information becomes electronic, the potential rewards to a dedicated commercial predator also increase significantly and may justify the commitment of adequate resources.

A serious effort — on the order of \$300,000 — by a legitimate or illegitimate business could find a DES key in an average of 19 days using off-the-shelf technology and in only 3 hours using a custom developed chip. In the latter case, it would cost \$38 to find each key (again assuming a 3 year life to the chip and continual use). A business or government willing to spend \$10,000,000 on custom chips, could recover DES keys in an average of 6 minutes, for the same \$38 per key.

At the very high end, an organization — presumably a government intelligence agency — willing to spend \$300,000,000 could recover DES keys in 12 seconds each! The investment required is large but not unheard of in the intelligence community. It is less than the cost of the Glomar Explorer, built to salvage a single Russian submarine, and far less than the cost of many spy satellites. Such an expense might be hard to justify in attacking a single target, but seems entirely appropriate against a cryptographic algorithm, like DES, enjoying extensive popularity around the world.

There is ample evidence of the danger presented by government intelligence agencies seeking to obtain information not only for military purposes but for commercial advantage. Congressional hearings in 1993 highlighted instances in which the French and Japanese governments spied on behalf of their countries' own businesses. Thus, having to protect commercial information against such threats is not a hypothetical proposition.

4.2 There are smarter avenues of attack than brute force.

It is easier to walk around a tree than climb up and down it. There is no need to break the window of a house to get in if the front door is unlocked.

Calculations regarding the strength of encryption against brute-force attack are *worst case* scenarios. They assume that the ciphers are in a sense perfect and that attempts to find shortcuts

have failed. One important point is that the crudest approach — searching through the keys — is entirely feasible against many widely used systems. Another is that the keylengths we discuss are always minimal. As discussed earlier, prudent designs might use keys twice or three times as long to provide a margin of safety.

4.3 The analysis for other algorithms is roughly comparable.

The above analysis has focused on the time and money required to find a key to decrypt information using the RC4 algorithm with a 40-bit key or the DES algorithm with its 56-bit key, but the results are not peculiar to these ciphers. Although each algorithm has its own particular characteristics, the effort required to find the keys of other ciphers is comparable. There may be some differences as the result of implementation procedures, but these do not materially affect the brute-force breakability of algorithms with roughly comparable key lengths.

Specifically, it has been suggested at times that differences in set-up procedures, such as the long key-setup process in RC4, result in some algorithms having effectively longer keys than others. For the purpose of our analysis, such factors appear to vary the effective key length by no more than about eight bits.

5 Appropriate Key Lengths for the Future — A Proposal

Table I summarizes the costs of carrying out brute-force attacks against symmetric cryptosystems with 40-bit and 56-bit keys using networks of general purpose computers, Field Programmable Gate Arrays, and special-purpose chips.

It shows that 56 bits provides a level of protection — about a year and a half — that would be adequate for many commercial purposes against an opponent prepared to invest \$10,000. Against an opponent prepared to invest \$300,000, the period of protection has dropped to the barest minimum of 19 days. Above this, the protection quickly declines to negligible. A very large, but easily imaginable, investment by an intelligence agency would clearly allow it to recover keys in real time.

What workfactor would be required for security today? For an opponent whose budget lay in the \$10 to 300 million range, the time required to search out keys in a 75-bit keyspace would be between 6 years and 70 days. Although the latter figure may seem comparable to the ‘barest minimum’ 19 days mentioned earlier, it represents — under our amortization assumptions — a cost of \$19 million and a recovery rate of only five keys a year. The victims of such an attack would have to be fat targets indeed.

Because many kinds of information must be kept confidential for long periods of time, assessment cannot be limited to the protection required today. Equally important, cryptosystems — especially if they are standards — often remain in use for years or even decades. DES, for example, has been in use for more than 20 years and will probably continue to be employed for several more. In particular, the lifetime of a cryptosystem is likely to exceed the lifetime of any individual product embodying it.

A rough estimate of the minimum strength required as a function of time can be obtained by applying an empirical rule, popularly called ‘Moore’s Law,’ which holds that the computing power available for a given cost doubles every 18 months. Taking into account both the lifetime of cryptographic equipment and the lifetime of the secrets it protects, we believe it is prudent to require that encrypted data should still be secure in 20 years. Moore’s Law thus predicts that the keys should be approximately 14 bits longer than required to protect against an attack today.

Bearing in mind that the additional computational costs of stronger encryption are modest, we strongly recommend a minimum key-length of 90 bits for symmetric cryptosystems.

It is instructive to compare this recommendation with both Federal Information Processing Standard 46, The Data Encryption Standard (DES), and Federal Information Processing Standard 185, The Escrowed Encryption Standard (EES). DES was proposed 21 years ago and used a 56-bit key. Applying Moore's Law and adding 14 bits, we see that the strength of DES when it was proposed in 1975 was comparable to that of a 70-bit system today. Furthermore, it was estimated at the time that DES was not strong enough and that keys could be recovered at a rate of one per day for an investment of about twenty-million dollars. Our 75-bit estimate today corresponds to 61 bits in 1975, enough to have moved the cost of key recovery just out of reach. The Escrowed Encryption Standard, while unacceptable to many potential users for other reasons, embodies a notion of appropriate key length that is similar to our own. It uses 80-bit keys, a number that lies between our figures of 75 and 90 bits.

Table I

| Type of Attacker | Budget | Tool | Time and cost per key recovered | | Length Needed for protection in Late 1995 |
|----------------------|----------|-------------------------------|---|--|---|
| | | | 40bits | 56bits | |
| Pedestrian Hacker | | | | | |
| | tiny | scavenged computer time | 1 week | infeasible | 45 |
| | \$400 | FPGA | 5 hours (\$0.08) | 38 years (\$5,000) | 50 |
| Small Business | | | | | |
| | \$10,000 | FPGA | 12 minutes (\$0.08) | 18 months (\$5,000) | 55 |
| Corporate Department | | | | | |
| | \$300K | FPGA or ASIC | 24 seconds (\$0.08) .18 seconds (\$0.001) | 19 days (\$5,000) 3 hours (\$38) | 60 |
| Big Company | | | | | |
| | \$10M | FPGA or ASIC | .7 seconds (\$0.08) .005 seconds (\$0.001) | 13 hours (\$5,000) 6 minutes (\$38) | 70 |
| Intelligence Agency | | | | | |
| | \$300M | ASIC | .0002 seconds (\$0.001) | 12 seconds (\$38) | 75 |

About the Authors

Matt Blaze is a senior research scientist at AT&T Research in the area of computer security and cryptography. Recently Blaze demonstrated weaknesses in the U.S. government's 'Clipper Chip' key escrow encryption system. His current interests include large-scale trust management and the applications of smartcards.

Whitfield Diffie is a distinguished Engineer at Sun Microsystems specializing in security. In 1976 Diffie and Martin Hellman created public key cryptography, which solved the problem of sending coded information between individuals with no prior relationship and is the basis for widespread encryption in the digital information age.

Ronald L. Rivest is a professor of computer science at the Massachusetts Institute of Technology, and is Associate Director of MIT's Laboratory for Computer Science. Rivest, together with Leonard Adleman and Adi Shamir, invented the RSA public-key cryptosystem that is used widely throughout industry. Ron Rivest is one of the founders of RSA Data Security Inc. and is the creator of variable key length symmetric key ciphers (e.g., RC4).

Bruce Schneier is president of Counterpane Systems, a consulting firm specializing in cryptography and computer security. Schneier writes and speaks frequently on computer security and privacy and is the author of a leading cryptography textbook, Applied Cryptography, and is the creator of the symmetric key cipher Blowfish.

Tsutomu Shimomura is a computational physicist employed by the San Diego Supercomputer Center who is an expert in designing software security tools. Last year, Shimomura was responsible for tracking down the computer outlaw Kevin Mitnick, who electronically stole and altered valuable electronic information around the country.

Eric Thompson heads AccessData Corporation's cryptanalytic team and is a frequent lecturer on applied cryptography. AccessData specializes in data recovery and decrypting information utilizing brute force as well as 'smarter' attacks. Regular clients include the FBI and other law enforcement agencies as well as corporations.

Michael Wiener is a cryptographic advisor at Bell-Northern Research where he focuses on cryptanalysis, security architectures, and public-key infrastructures. His influential 1993 paper, Efficient DES Key Search, describes in detail how to construct a machine to brute force crack DES coded information (and provides cost estimates as well).

ACKNOWLEDGEMENT

The authors would like to thank the Business Software Alliance, which provided support for a one-day meeting, held in Chicago on 20 November 1995.