

Security Estimates for 512-bit RSA

M.J.B. Robshaw
RSA Laboratories
100 Marine Parkway
Redwood City, CA 94065-1031
`matt@rsa.com`

June 29, 1995

Abstract

In this note we address the short-term security offered by the use of a 512-bit RSA modulus. Following recent tremendous improvements to the practicality of the generalized number field sieve, it must be expected that by the end of next year, a 512-bit RSA number will have been factored. However, for those fielded systems which use 512-bit RSA, what are the implications? Some systems may well continue using 512-bit RSA long after one particular 512-bit RSA number has been factored. In this note, we present data which might provide answers to questions about the continuing use of a 512-bit RSA modulus.

1 Introduction

It is well known that the security of the RSA cryptosystem [4] relies on the continuing impracticality of factoring large numbers of a particular type. It is also well known that advances in factoring techniques, together with continual improvements in hardware performance, mean that increasingly large and ‘difficult’ numbers can be factored as time goes by.

Using a variety of techniques it is possible to estimate, with reasonable accuracy in the short term, the size of the modulus that should be used in an implementation of RSA to attain some desired security level. However, there are few estimates which provide information on the increasing vulnerability of systems with a specific, and perhaps fixed, size of modulus.

For some considerable time a 512-bit RSA modulus has been considered as offering relatively good security. However, recent improvements in factoring techniques force us to closely consider the increasing vulnerability of a 512-bit RSA modulus. In this note we try to estimate the likely risks involved in continuing to use a 512-bit modulus over the next ten years.

2 Three projections

Throughout this note we will make some basic assumptions. Foremost among them is that when using the generalized number field sieve [2] (which is at present the most effective algorithm for tackling larger RSA numbers) then the number of MIPS-years¹ (abbreviated to MY) required to factor a 512-bit RSA modulus is roughly 3×10^4 [3].

Additionally we will assume that the computing power per dollar doubles every 18 months (a common assumption) and that a 10 MIPS machine (or the parts thereof) can currently be bought for U.S.\$500.

In this note we have taken no account of potential algorithmic improvement. However, it is worth noting that a number with special form and a length of 162 decimal digits was recently factored using the *special number field sieve* [3]. While a number of 162 decimal digits is longer than one of 512 bits, this factorization required a particularly modest 200 MY. If anywhere near this kind of algorithm performance can be delivered on numbers without this special form, then 512-bit RSA numbers will be truly weak.

2.1 Dollar investment

In this section we consider adversaries who are buying and setting up equipment exclusively to factor 512-bit RSA numbers. The estimates that follow can easily

¹The number of operations completed in a year by a machine operating at one million instructions per second.

be adjusted if it is felt that the cost of buying powerful computing equipment is quite different from our presumed figure in 1995 of 10 MIPS for U.S.\$500.

In the table that follows, we estimate the time in years required to factor a 512-bit RSA number with an investment of the dollar amount shown. We assume that the increase in computing power translates into a drop in purchasing cost. Note that in this table and others following, some numbers appear to remain unchanged between successive years due to rounding.

	<i>year</i>				
<i>Investment</i>	<i>1996</i>	<i>1997</i>	<i>1998</i>	<i>1999</i>	<i>2000</i>
\$100,000	9.5	6.0	3.8	2.4	1.5
\$1,000,000	0.9	0.6	0.4	0.2	0.2
\$10,000,000	0.1	<	<	<	<
	<i>2001</i>	<i>2002</i>	<i>2003</i>	<i>2004</i>	<i>2005</i>
\$100,000	0.9	0.6	0.4	0.2	0.2
\$1,000,000	0.1	<	<	<	<
\$10,000,000	<	<	<	<	<

Table 1: The estimated time in years required to factor a 512-bit RSA number with a given investment of 1996–2005 technology. The symbol < is used to denote less than one month.

2.2 Well positioned adversary

In this section we consider the role of a systems administrator or some other individual with access to a considerable amount of computing power within a company. Such an individual might obtain factoring software from the Internet and use the spare cycles on company machines to attack the keys used by other employees. This consideration was motivated by calculations performed by Odlyzko [3].

While an individual might attack one of the workers’ keys chosen at random; the key used by a finance or policy director would most likely be the key under threat.

In the following table we give the time required in years to attack a 512-bit RSA modulus. We assume that each machine in the company has an effective rating of 10 MIPS in 1995 (since spare cycles are being used in the illicit factorization). Machines that are considerably faster may well be available and would alter the predictions accordingly. The number of workstations were chosen to represent a small company, a large company and a very large company.

	<i>year</i>				
<i># workstations</i>	<i>1996</i>	<i>1997</i>	<i>1998</i>	<i>1999</i>	<i>2000</i>
50	38	24	15	9.5	6.0
500	3.8	2.4	1.5	0.9	0.6
5,000	0.4	0.2	0.2	0.1	<
	<i>2001</i>	<i>2002</i>	<i>2003</i>	<i>2004</i>	<i>2005</i>
50	3.8	2.4	1.5	0.9	0.6
500	0.4	0.2	0.2	0.1	<
5,000	<	<	<	<	<

Table 2: The estimated time in years required to factor a 512-bit RSA number, with a given number of workstations, for the years 1996–2005. The symbol < is used to denote less than one month.

2.3 Network attack

The increasing use of the Internet to network together computers is an important feature in contemporary factoring efforts. In this section we try to estimate the number of people required to participate in some widely publicized factoring effort.

We shall assume that each person has access, and is willing to offer, 10 MIPS worth of processing power in 1995 and an increasing amount in successive years in line with improving hardware performance. In this third table we present the number of people (or workstations) which are needed to factor a single 512-bit RSA modulus in the time shown.

	<i>year</i>				
<i>time</i>	<i>1996</i>	<i>1997</i>	<i>1998</i>	<i>1999</i>	<i>2000</i>
<i>2 years</i>	950	600	380	240	150
<i>1 year</i>	1900	1200	750	470	300
<i>6 months</i>	3800	2400	1500	950	600
	<i>2001</i>	<i>2002</i>	<i>2003</i>	<i>2004</i>	<i>2005</i>
<i>2 years</i>	94	59	37	23	15
<i>1 year</i>	190	120	74	47	30
<i>6 months</i>	380	240	150	94	59

Table 3: The number of people that might be required to collaborate on achieving the factorization of a 512-bit RSA number in a given time period, for the years 1996–2005.

3 New techniques

In our analysis we have made no allowance for any potential algorithmic improvement. How much might this undermine the figures we have presented?

We have restricted our attention to a specific 10-year span, and it is unclear how much algorithmic improvement should be allowed for within this period. Perhaps there will be no substantial improvements within the next 10 years, or perhaps any improvements will only be significant when applied to numbers much longer than 512 bits.

It is unlikely, however, that this will in fact be the case. Already it is known that the work effort required in April 1994 for the landmark factorization of RSA-129 [1] could now be considerably reduced using the newer generalized number field sieve.

We must stress that significant improvements to the number field sieve will dramatically undermine any security offered by 512-bit moduli. Even if the best improvements to the generalized number field were to make it only three times more efficient than today², then

- \$1,000,000 could be spent *next year* to factor a 512-bit modulus in under four months rather than almost a year,
- a company with 5,000 workstations *next year* would have sufficient resources to factor a 512-bit modulus in under six weeks rather than around five months, and
- a team of fewer than twice the number of people involved in the factorization of RSA-129 could factor a 512-bit modulus in about six months *next year*.

These should be serious considerations.

4 Conclusions

The figures in this report represent important implications for the use of 512-bit RSA moduli, even in the short term.

Considering the effort required to factor a 512-bit RSA modulus, we have observed that for an investment in 1997 of \$1,000,000, an attacker might be expected to factor a 512-bit modulus in under eight months. With a \$10,000,000 investment in 1996, it might take around five weeks. While such an investment is perhaps out of proportion to the value of the data protected using a 512-bit modulus, the investment can be recovered by factoring other numbers. And, as we have repeatedly stressed, we have taken no account of future algorithm improvements. It is not inconceivable that by the turn of the century, a viable

²This would make the effort required to factor a 512-bit number around 10^4 MY.

business could be established that is dedicated to factoring 512-bit numbers, assuming there is a market for such an enterprise.

The power of the well positioned adversary should demand the utmost attention. As we have shown, the computing power *already* possessed in 1995 by large companies³, could be harnessed to factor 512-bit numbers in a matter of months.

Perhaps less significant is the short-term risk posed by overt networked factoring efforts. Since 512-bit keys should not be used to secure any valuable information it is unlikely that nearly a thousand people could be persuaded to donate spare cycles to factor a 512-bit RSA number in a year. (Within a few years we might assume that the novelty of factoring 512-bit numbers has worn off.) However, networked attacks are an important consideration for certification hierarchy root-keys which are high-profile targets. Clearly, the moduli for such valuable data should be chosen to be well out of reach of even the most committed efforts.

With the current information we have, it might be reasonably argued that with regards to the installed base of 512-bit RSA, it will still be moderately expensive to attack any individual key over the next two or three years. Indeed, if improvements to the speed of computation provide the only advances in factoring ability in the near future, then the figures in this note might be used to give a rough idea of the increasing risk incurred by the continued use of a 512-bit RSA key.

It cannot be sufficiently stressed however, that any new advances in the performance of factoring algorithms will, in all likelihood, have catastrophic implications for the security offered by 512-bit RSA numbers.

Predictions for larger RSA moduli are notoriously difficult to make; new developments often overtake old predictions. However, recent estimates [3] put the computational effort required to factor a 768-bit RSA modulus, using the current techniques that threaten 512-bit RSA, at 2×10^8 MY and to factor a 1024-bit RSA modulus, at 3×10^{11} MY. These contrast with the 3×10^4 MY estimated to attack a 512-bit modulus and they clearly offer a much more acceptable level of security.

For many years now, 512-bit RSA has been adequate for a great number of applications. At present however, even without further advances in factoring techniques, 512-bit RSA can only be considered as offering moderate, short-term security.

The first conclusion to draw from this note is that systems implementing RSA should support a variable key size. While it is easy to account for improvements in computational performance, one should also allow for unforeseen developments in factoring ability when choosing the size of an RSA modulus.

And perhaps most importantly, it is clear that factoring 512-bit RSA moduli

³Odlyzko estimates that Silicon Graphics has around 10,000 workstations which could each contribute 10 MIPS.

is soon going to be moderately routine. With this in mind, users are advised to begin phasing out the use of 512-bit RSA as soon as possible. Depending on the particular security requirements, and assuming that there are no new developments in factoring ability, some users might prefer to continue using 512-bit RSA moduli for moderate to low security applications. Even for this use, however, it seems prudent to recommend that 512-bit RSA moduli should not be used after 1997 or 1998 at the latest.

References

- [1] D. Atkins, M. Graff, A.J. Lenstra, and P.C. Leyland. The magic words are squeamish ossifrage. In J. Pieprzyk and R. Safavi-Naini, editors, *Advances in Cryptology — Asiacrypt '94*, pages 263–277, Springer-Verlag, Berlin, 1995.
- [2] J.P. Buhler, H.W. Lenstra, and C. Pomerance. Factoring integers with the number field sieve. 1992. To appear.
- [3] A. Odlyzko. The future of integer factorization.⁴ *CryptoBytes*, 1(2), Summer 1995. To appear.
- [4] R.L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.

⁴Also available by sending the message *future.of.factoring.ps* from *att/math/odlyzko* to *netlib@research.att.com*.