# ETSI GR QSC 006 V1.1.1 (2017-02)


# Quantum-Safe Cryptography (QSC); Limits to Quantum Computing applied to symmetric key sizes

***Disclaimer***

The present document has been produced and approved by the Quantum-Safe Cryptography (QSC) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference
DGR/QSC-006

Keywords
cyber security, quantum cryptography, security

***ETSI***

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

# Contents

# Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

# Foreword

This Group Report (GR) has been produced by ETSI Industry Specification Group (ISG) Quantum-Safe Cryptography (QSC).

# Modal verbs terminology

In the present document "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

# Executive summary

The present document analyses the impact of a quantum computer on symmetric cryptographic primitives. A worst-case estimate is derived for the maximum available quantum computing power in 2050. This leads to the conclusion that 256-bit symmetric ciphers and hash functions will still be unbroken in 2050.

# Introduction

A quantum computer will require an enormous change in the cryptographic landscape [i.7]. This is why research and standardization effort is put into finding quantum-safe asymmetric alternatives for RSA, (EC) Diffie-Hellman, and (EC)DSA. Significant effort from industry will be put into preparing for the necessary transition to these new asymmetric primitives.

However, symmetric primitives like AES, SHA-2, and SHA-3 are equally integrated into the numerous information security solutions that exist worldwide. Since a quantum computer can also speed up attacks on symmetric primitives [i.6], it is important to analyse how long these symmetric primitives - and their most-used key sizes - will remain secure.

The present document studies the long-term security of symmetric primitives such as AES-256, SHA-2, and SHA-3. A scientific approach shows that attacks cannot continue to improve at an exponential rate forever. Moore's Law may assert that transistors become twice as small roughly every 1,5 years, but this trend cannot continue and in fact has already stopped. While it is unknown whether a similar trend will appear for quantum computers, it is possible to put an upper bound on the quantum computing power that could be developed in the foreseeable future. The analysis in the present document is based on conservative assumptions and estimates. This does not result in exact dates on when each primitive will be broken, but it does assert their security for at least a certain period of time.

The present document concludes that there are existing and widely used symmetric (AES-256) and hash primitives (SHA-2 and SHA-3 with an output length of at least 256 bits) that will withstand quantum computer attacks until way after 2050. It is reassuring to know that for these symmetric primitives there is no need to find and heavily scrutinize alternatives within the next few years, like is done for the asymmetric primitives.

Note that this does not mean that there is no need to look into symmetric algorithms when it comes to the threat of a quantum computer. On the contrary, industry does have to worry about symmetric algorithms, since there are billions of devices in the world that rely on a symmetric cipher with a key length of 128 bits or less. Examples include mobile communication with e.g. GSM or TETRA. Unfortunately, the calculations that are used in the present document to assert that AES-256 will remain secure until way after 2050 cannot be used to predict when a quantum computer can attack AES-128, or any other cipher with a short key length. Therefore, industry is advised to identify where their products rely on smaller key and hash output lengths, and to start investigating the necessary steps for a transition to primitives with key lengths that will withstand quantum computer attacks like the ones investigated in the present document.

# 1 Scope

The present document gives information on the long-term suitability of symmetric cryptographic primitives in the face of quantum computing.

# 2 References

## 2.1 Normative references

Normative references are not applicable in the present document.

## 2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] AI Impacts (March 2015): "Trends in the cost of computing".

NOTE: Available at http://www.aiimpacts.org/trends-in-the-cost-of-computing.

[i.2] Thomas Monz et al (2011): "14-Qubit Entanglement: Creation and Coherence", Phys. Rev. Lett. 106, 130506.

[i.3] Christof Zalka: "Grover's quantum searching algorithm is optimal", Phys. Rev. A 60, 2746, 1999, arXiv.

NOTE: Available at http://www.arxiv.org/abs/quant-ph/9711070.

[i.4] PriceWaterhouseCoopers, The world in 2050 (February 2015): "Will the shift in global economic power continue?".

NOTE: Available at www.pwc.com/gx/en/issues/the-economy/assets/world-in-2050-february-2015.pdf.

[i.5] World Bank, Data: "Research and development expenditure" (% of GDP).

NOTE: Available at http://data.worldbank.org/indicator/GB.XPD.RSDV.GD.ZS.

[i.6] Lov K. Grover: "A fast quantum mechanical algorithm for database search", STOC 1996, pp 212-219, ACM 1996.

[i.7] Peter W. Shor: "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer", SIAM Journal on Computing, 26(5):1484-1509, 1997.

[i.8] Markus Grassl, Brandon Langenberg, Martin Roetteler, and Rainer Steinwandt (December 2015): "Applying Grover's Algorithm to AES: quantum resource estimates".

[i.9] Matthew Amy, Olivia Di Matteo, Vlad Gheorghiu, Michele Mosca, Alex Parent, and John Schanck (March 2016): "Estimating the cost of generic quantum pre-image attacks on SHA-2 and SHA3".

[i.10] Marc Kaplan, Gactan Leurent, Anthony Leverrier, and María Naya-Plasencia (February 2016): "Breaking symmetric cryptosystems using quantum period finding".

[i.11] Andrey Bogdanov, Dmitry Khovratovich, and Christian Rechberger (2013): "Biclique analysis of the full AES".

[i.12] Daniel J. Bernstein (May 2009): "Cost analysis of hash collisions: will quantum computers make SHARCS obsolete?".

[i.13] Simon J. Devitt, William J. Munro, and Kae Nemoto (June 2013): "Quantum Error Correction for Beginners", Rep. Prog. Phys. 76 (2013) 076001, arXiv:0905.2794.

[i.14] European Commission D-G for Research and Innovation: "Global Europe 2050", European Union 2012, DOI: 10.2777/79992.

[i.15] Arjen K. Lenstra and Eric R. Verheul (2001): "Selecting Cryptographic Key Sizes", Journal of Cryptology 14, 4, pp 255-293, Springer Berlin Heidelberg.

[i.16] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland (September 2012): "Surface codes: Towards practical large-scale quantum computation", Phys. Rev. A 86, 032324.

# 3 Symbols and abbreviations

## 3.1 Symbols

For the purposes of the present document, the following symbols apply:

| | |
|---|---|
| $ | US Dollar |
| d | days |
| EHz | exahertz |
| h | hours |
| Hz | hertz |
| nm | nanometre |
| PHz | petahertz |
| pm | picometre |
| y | years |

## 3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

| | |
|---|---|
| AES | Advanced Encryption Standard |
| DSA | Digital Signature Algorithm |
| EC | Elliptic Curve |
| EU | European Union |
| GDP | Gross Domestic Product |
| GSM | Global System for Mobile communications |
| MAC | Message Authentication Code |
| MIPS | Million Iterations Per Second |
| QC | Quantum Computer |
| QEC | Quantum Error Correction |
| RSA | Rivest Shamir Adleman |
| SHA | Secure Hash Algorithm |
| TETRA | Terrestrial Trunked Radio |
| USA | United States of America |

# 4 Background

## 4.1 Asymmetric cryptography and quantum computing

If a large quantum computer is built, it would pose a threat to several cryptographic primitives. Most notably, RSA, (EC) Diffie-Hellman, and (EC)DSA would be completely broken by Shor's algorithm [i.7]. Here completely broken means that, given complete quantum control over a sufficient number of qubits, a reasonable size ($O(n^3)$) quantum circuit can break the underlying mathematical problem in reasonable ($O(n^3)$) time for a private key of $O(n)$ bits. This ignores the fact that the cryptographic primitive needs to be implemented in a quantum circuit and also ignores *quantum error correction* (QEC) to avoid decoherence. Both would introduce a polynomial overhead. QEC is a technique that allows stable computation with unstable qubits. Since reading out qubits destroys their quantum properties this technique is more involved than normal error correction techniques. QEC introduces a significant overhead in the circuit size and computation time [i.8], [i.9]. Most types of qubits tend to be unstable so it seems likely that QEC will be needed in a large quantum computer. For a more in-depth treatment of QEC see [i.13].

## 4.2 Symmetric cryptography and quantum computers

For symmetric algorithms the impact of quantum computers is less clear. A rule of thumb says that Grover's algorithm effectively halves the key size for these algorithms [i.6]. However, the aforementioned 'polynomial overhead' considerably increases the complexity of this algorithm: breaking a 128-bit AES key costs about $2^{87}$ gates and takes the time of $2^{81}$ gate operations [i.8] rather than $2^{64}$ operations predicted by the rule of thumb. Finding pre-images for SHA-2 and SHA-3 also has a considerable overhead, costing the time equivalent to $2^{166}$ hash function calls for both SHA-2 and SHA-3 [i.9], where the rule of thumb would predict $2^{128}$. The footprint of QEC further increases the number of qubits from a few thousand to more than 10 million [i.9]. Since [i.8] attempts to minimize the number of qubits, while [i.9] attempts to minimize the T-gate depth, different qubit and operation counts can be obtained for different implementations.

Existing symmetric algorithms might be vulnerable to other quantum attacks. For example [i.10] demonstrates that several MAC and authenticated encryption modes can be broken with a quantum computer if an attacker has access to a *quantum implementation* of the primitive and can query it with *superpositions*, which seems quite a strong assumption. The present document assumes the use of algorithms that have no structural weaknesses that can be exploited by a (quantum) adversary. This means that breaking such an algorithm is the same as solving the general search problem.

Grover's algorithm is optimal for solving the general search problem. It solves the general search problem on a set of size $N$ in $O(\sqrt{N})$ time, while no quantum algorithm exists that solves this problem faster [i.3]. In addition, implementing Grover's algorithm in parallel results in a classical time-memory trade-off: $m$ quantum computers can solve the general search problem in no less than $O(\sqrt{N/m})$ time, which can trivially be achieved by partitioning the problem into $m$ problems of size $N/m$. The total cost of this parallel computation is $O(\sqrt{Nm})$ so it is more efficient not to parallelise the computation. As an example, one quantum computer could find a 256-bit AES key in about $2^{128}$ time, while $2^{32}$ parallel quantum computers could find this key in about $2^{112}$ time. The overall cost of the latter computation is about $2^{144}$ which is much more than the $2^{128}$ cost of the single quantum computer. Nevertheless, parallelisation might still be a sensible choice because no adversary is willing to wait longer than a few years for a decryption.

## 4.3 Number of qubits

As stated before, possibly millions of physical qubits are needed to break a 256-bit symmetric key. In classical computers, the amount of available memory also follows a version of Moore's Law. For the number of qubits a different behaviour is expected.

While qubits can be built, the main challenge is creating qubits that are both stable, and can be used for quantum computations. The world record of 14 entangled qubits was set in 2011 [i.2]. Significant progress in this area is not expected until a stable logical qubit is created. This stable qubit could be constructed from multiple physical qubits using error correction, or it could be inherently physically stable. Once this is achieved, the number of qubits can grow very rapidly. There is no reason to expect that this growth is limited by Moore's Law, or at least not until millions or billions of qubits have been reached. Therefore, once stable qubits are available, the number of qubits is not a limiting factor for any cryptographic attacks.

## 4.4 Outline of the present document

Assuming a symmetric algorithm is used without structural weaknesses and that the algorithm is not implemented as a quantum random oracle, can it be broken by a quantum computer? What key lengths are safe to use? What if qubits become more stable and QEC techniques are improved? What if Moore's Law applies to quantum computers? The present document positively answers these questions based on worst-case assumptions. Two scenarios are analysed in clause 5: this clause estimates the capabilities of the fastest quantum computer that could ever be built, and it gives an optimistic estimate of the commercially available quantum computing power around 2050. Clause 6 analyses which key sizes are still safe in 2050, and clause 7 gives the conclusions. All examples will focus on finding a 256-bit symmetric key, but general expressions will be derived to address general key sizes of *k* bits.

# 5 Quantum computers in 2050

## 5.1 Approach

This clause analyses two scenarios for quantum computers: the *commercial quantum computer* and the *worst case quantum computer*. The *commercial quantum computer* is a very optimistic estimate of a quantum computer that could be commercially available in 2050. It is much faster than today's regular computers, costs about the same, and needs only a single clock cycle per Grover iteration. Its qubits are arranged on a flat surface and its computations are kept stable with QEC. The *worst case quantum computer* is a special-purpose extremely optimized quantum computer where the qubits are inherently stable and QEC is not needed. Its qubits are packed together as closely as possible to maximize the clock speed.

For each type of quantum computer, its cost is estimated by a version of Moore's Law. Assuming that an attacker has a limited budget this gives an upper bound for an attacker's computing power. This upper bound is used in clause 6 to derive which symmetric key sizes are safe for the foreseeable future.

## 5.2 Moore's Law

Moore's Law is a well-recognized trend that enables making rough estimates of future computing power. Current technology is reaching the limits of Moore's Law. Some further miniaturization may be possible through clever engineering, but transistors cannot be made smaller than the size of a single atom. Even that size will not be reached, because transistors rely on bulk behaviour, which is only possible in systems that are significantly larger than the single atom scale.

Vendors are shipping their products with an increasing number of processing cores to keep up the illusion of Moore's Law. Due to cheaper production techniques this can be done without a significant increase in product price. So the processing power per dollar still manages to follow Moore's Law. According to [i.1] the processing power per dollar follows the following equation:

$$\text{MIPS/\$ (year)} = 10^{-360{,}109288 + \text{year} * 0{,}178929} \quad (1)$$

In other words: Moore's Law will cease to be an accurate measure of processing power per square millimetre of silicon but may remain relevant when estimating the processing capability per dollar in the future. Since Grover's algorithm does not parallelise well, clock speeds are also relevant. Commercial chips have increased their clock speed for decades, but are now stagnating around a maximum of about 4 GHz. THz clock speeds are achievable, but these are hard to realize in practice because they require more heat dissipation. Therefore, these speeds do not seem interesting for commercial applications. Increasing clock speeds even further, beyond the THz regime, is currently inconceivable but not fundamentally impossible.

## 5.3 'Commercial' quantum computers

The next two clauses derive some fundamental limits for quantum computers. This clause focuses on 'commercial' quantum computers, while the next clause analyses worst case quantum computers. It uses a conservative approach and emphasizes any assumptions that directly affect the conclusions. To get some numbers to work with, the following assumptions are made:

- **Assumption 1:** qubits will not become smaller than a single atom.
- **Assumption 2:** information between qubits cannot be exchanged faster than the speed of light.

The first assumption implies a set of qubits tightly packed together with quantum circuits that connect them. There are multiple technologies for building a qubit. While a qubit itself could technically be smaller than a single atom, it seems impossible to include control logic at the subatomic level and still avoid decoherence. For the second assumption note that while entanglement is a powerful resource for quantum computation, it cannot be used to convey information faster than the speed of light. In today's electronic circuits, information is communicated with a speed that is in the same order of magnitude as the speed of light.

The sequel focuses on a quantum computer with at least $k$+1 = 257 logical qubits, because this is theoretically the minimum number needed to find a 256 bit key with Grover's algorithm. With quantum error correction at least a few million physical qubits are needed [i.9], and possibly a lot more; for AES, SHA-2 and SHA-3, a few thousand logical qubits are needed, resulting in about $10^7$ physical qubits [i.8], [i.9]. Since [i.8] attempts to minimize the number of qubits, minimizing the computation time could require even more qubits. For a detailed analysis of the overhead of quantum error correction and scalability problems due to qubit wiring and classical control electronics see [i.16]. The assumption is that due to a breakthrough in QEC, each logical qubit in a commercial quantum computer is made of no more than 100 atoms:

- **Assumption 3:** a logical qubit requires at least 100 atoms.

Single atom qubits will be addressed in clause 5.4 which assesses the worst case quantum computer. For a commercial quantum computer the qubits are assumed to be made of silicon, which has a Van der Waals radius of 210 pm. Most materials actually have a larger Van der Waals radius. Atoms can be compressed slightly below their Van der Waals radius if they form chemical bonds, but this is neglected for now. When the 25 700 qubits are packed together as closely as possible on a 2-dimensional surface, the result is a circle with an area of 3 600 nm$^2$. The diameter $d$ of this circle is given by:

$$d = 4{,}2\sqrt{k+1}\ \text{nm}, \tag{2}$$

so k = 256 gives $d$ = 67 nm. Communicating from one side to the other side of the circle takes about $2{,}1 \cdot 10^{-16}$ seconds, meaning that a quantum computer with this number of qubits cannot run at a frequency higher than $4{,}5 \cdot 10^{15}$ Hz or 4,5 PHz. In general, the maximum frequency of such a quantum computer is given by:

$$f = \frac{7{,}1 \cdot 10^{16}}{\sqrt{k+1}}\text{Hz}. \tag{3}$$

This analysis assumes that this is actually the number of qubits operations per second, even though fault-tolerant operation with QEC may require more than one operation per qubit instruction. Three-dimensional configurations seem hard to achieve, since cooling and circuitry to connect the qubits are needed. So this seems the best achievable case for commercial applications.

## 5.4 Worst case quantum computers

For special purpose quantum computers, an absolute worst case estimate of the maximum frequency is derived. The analysis is similar to the previous clause, but uses different assumptions. It assumes each qubit, including control logic, is only the size of a single atom. In addition, it assumes the qubits to be inherently stable so that QEC is not needed. To make this quantum computer as fast as possible, the smallest possible atoms (hydrogen) are packed together as closely as possible in a spherical packing. In short: Assumptions 1 and 2 are kept, while replacing Assumption 3 by the following:

- **Assumption 3':** a logical qubit is as large as a single atom and is inherently stable.

Hydrogen has a Van der Waals radius of only 120 pm. Without the need of QEC $k$+1 qubits suffice. In this case the distance $d = 2r$ between the qubits that are furthest apart can be derived from $\frac{4}{3}\pi r^3 = (k+1)\frac{4}{3}\pi(120)^3$ and is given by:

$$d = 240\sqrt[3]{k+1}\text{ pm.} \quad (4)$$

For $k$ = 256 this gives $d$ =1,5 nm. Light takes $d/c = 5 \cdot 10^{-18}$ seconds to travel this distance so the maximum frequency is $2{,}0 \cdot 10^{17}$ Hz = 0,2 EHz. In general the maximum frequency is given by:

$$f = \frac{1{,}25 \cdot 10^{18}}{\sqrt[3]{k+1}}\text{ Hz.} \quad (5)$$

Higher speeds could be achieved if direct communication between each pair of qubits is not required. Cryptographic algorithms such as AES are designed to ensure that each key bit influences each ciphertext bit after as few rounds as possible. So when searching for an AES key Grover's algorithm will require interaction between each pair of qubits. If only the direct communication between adjacent qubits (240 pm) is considered, the maximum frequency would be 1,3 EHz. This value of 1,3 EHz is used as a conservative upper bound.

## 5.5 An upper bound for quantum computing budgets

The previous clauses put upper bounds on the speed of a quantum computer. But assessing which key sizes are secure also requires a bound on how many quantum computers could be used to break a key. This clause derives such a bound.

It is hard to predict how many worst case quantum computers with EHz speeds could be built and how much they would cost. Note that this would be a very ambitious and costly project and that one might be better off to wait for quantum computers to become commercially available. In this case, quantum clusters for parallel computing will become affordable. It follows from the assumptions in clauses 5.1 and 5.2 that commercial quantum processors achieve 4,5 PHz around 2050 and cost $900 a piece. Equation (1) gives $5 \cdot 10^6$ MIPS/$ in 2050, which results in $908 for a normal computer chip that runs at that speed.

Looking at the expected top global economies in 2050 [i.14], [i.4] and their research and development expenditure [i.5], suggests an upper bound for the maximum amount of money an attacker would be willing to spend on quantum computing. On average about 2 % of GDP is spent on research and development with peaks up to 4,1 % in some countries. The largest economy in 2050 (China) is projected at a GDP of $34 to $61 trillion. The next largest are the USA ($22 to $41 trillion), India ($11 to $42 trillion) and the EU ($20 to $26 trillion). The suggested upper bound assumes a large economy ($61 trillion) spends 4,1 % of its GDP on quantum computing.

This bounds the budget to at most $2,5 trillion which results in a cluster of $2{,}8 \cdot 10^9$ quantum processors running in parallel, each of which runs at 4,5 PHz. This cluster is as powerful as about 37 000 quantum computers running at 1,3 EHz. Note that a budget of $2,5 trillion may allow one to build a small cluster of EHz quantum computers, but a cluster of more than 37 000 seems far-fetched. The technology involved is so extreme that a cluster of commercially available quantum computers will probably be more cost effective.

Note, again, that $2,5 trillion is an extremely conservative upper bound, not an actual projected budget for quantum computing. No country will spend its entire research and development budget on quantum computing. For comparison: $2{,}8 \cdot 10^9$ quantum computers in 2050 is one quantum computer for every three people in the world, and is roughly equal to the total number of personal computers in the world in 2016.

# 6 Key and parameter sizes

## 6.1 Approach

This clause combines the results from clause 5 to assess which symmetric key sizes will remain secure for the foreseeable future. The analysis in clause 6.2 focuses on the security of AES-256 in 2050. Clause 6.3 analyses the long-term security of SHA-2 and SHA-3.

## 6.2 Symmetric keys

Table 6.1 compares the capabilities of a cluster of $2{,}8 \cdot 10^9$ quantum computers running at 4,5 PHz (see clause 5.3) to a single quantum computer running at 1,3 EHz. The time *t* to break a symmetric key of *n* bits has been calculated as:

$$t = \frac{2^{n/2}}{f \cdot \sqrt{\#QC}} \text{ seconds,} \quad (6)$$

where *#QC* is the number of quantum computers and *f* is the clock frequency of each quantum computer. Note that this approach assumes that each step of the computation takes only a single operation, so in practice it will take longer to break these keys.

**Table 6.1: Estimated time to break a symmetric key with a quantum computer (cluster) in 2050**

| Key size | Single 'worst case' QC at 1,3 EHz | $2{,}8 \cdot 10^9$ 'commercial' QC's at 4,5 PHz |
|---|---|---|
| 160 bits | 11 d | 1,4 h |
| 180 bits | 30 y | 61 d |
| 200 bits | 30 921 y | 170 y |
| 220 bits | - | 174 000 y |

It follows from Equation (6) that a 256-bit symmetric key could be broken in about 50 billion years, which means *one* AES-256 key could be broken based on worst-case assumptions and a 50-billion-year effort starting around 2050. This does not take into account possible cryptanalytic progress. The model for cryptanalytic progress from [i.15] assumes that attacks become twice as effective every 18 months. While not every cryptanalytic breakthrough implies a speed-up in Grover's algorithm applied on the corresponding primitive, a conservative assumption is that quantum attacks become $\sqrt{2}$ times as effective every 18 months, or twice as effective every 36 months. In general, the effective key size $n_{eff}$ in bits of a symmetric algorithm is given by:

$$n_{eff} = n - \frac{y - y_{ref}}{1{,}5}, \quad (7)$$

where *n* is the algorithm's key length, *y* is the current year, and $y_{ref}$ is the year when the algorithm was still considered full strength. Of course, this equation only works if *y* is not smaller than $y_{ref}$. To accommodate for cryptanalytic progress, simply replace *n* by $n_{eff}$ in Equation (7).

Biclique attacks on AES were published in 2013 [i.11], reducing effective key sizes by 2 to 4 bits. So as a reference, assume AES-*n* has its full effective key size of *n* bits in $y_{ref} = 2010$. This means that the effective key size of AES-256 would be 229 bits in 2050, which could be broken in about 4 million years using the largest quantum computer cluster imaginable. This does not imply that the encrypted data will remain secure for another 4 million years, but it does mean that this data will remain secure until at least 2050. If such cryptanalytic progress is actually made on AES, it will probably be deprecated and a different algorithm will be standardized. However, standardizing a new algorithm will not protect data that is encrypted and transmitted today. Therefore, it is reassuring that even in this scenario AES-256 will keep the encrypted data confidential until at least 2050.

## 6.3 Hash output lengths

Finding a pre-image for a hash function with output length *n* with Grover's algorithm is approximately as hard as finding a symmetric key of the same length. Therefore the crude estimates from the present document give the same results for both. This means that finding a pre-image for SHA-256 would take approximately 4 million years on a large quantum computer cluster, assuming cryptanalytic progress is made on SHA-256.

For hash function collisions, using a quantum computer does not result in a speedup [i.12]. The best approach would be to use a parallel version of Pollard's rho method on a classical computer cluster. The time needed to find hash collisions can be estimated from Equation (1). For example, finding a collision for SHA-256 requires about $2^{128}$ hash function compressions. This costs about $2^{136}$ instructions on today's hardware, but to accommodate for improvements in hardware assume it takes $2^{128}$ instructions in 2050. In 2050 computing power costs about $5 \cdot 10^6$ MIPS/$, so an attacker with a $2,5 trillion budget could have $10^{19}$ MIPS of computing power. Finding a collision for SHA-256 would then take about 1 million years. Taking cryptanalytic progress into account as before gives 75 years instead, which is less than the time required to break an AES-256 key on a quantum computer derived in clause 6.1. This is surprising because both have the same complexity of $2^{128}$. The reason is that parallelising Grover's algorithm is relatively inefficient.

The analysis for SHA-3 is the same as for SHA-2, so finding a pre-image for SHA3-256 on a large quantum computer cluster would also take approximately 4 million years. So the analysis shows that all versions of SHA-2 and SHA-3 with an output length of 256 bits or more will remain secure until at least 2050.

# 7 Conclusions

The present document analyses the impact of a quantum computer on symmetric cryptography. This analysis is based on conservative assumptions and estimates of increased computing power of future quantum computers, cryptanalytic progress, and budgets. The conclusion is that:

1) **It seems unlikely that a quantum computer will ever be faster than 1,3 EHz (worst case)** assuming that:
   - qubits cannot be packed more closely than atoms;
   - exchange of information between qubits is done at the speed of light.
2) **It seems unlikely that an adversary in 2050 can afford more than either 2,8 billion 'commercial' (4,5 PHz) quantum computers or 37 thousand 'worst case' (1,3 EHz) quantum computers** assuming:
   - Moore's Law;
   - estimated largest R&D budget of $2,5 trillion in 2050;
   - 'commercial' quantum computers will be as expensive as normal computers in 2050;
   - building a 1,3 EHz quantum processor will cost billions.
3) **It seems unlikely that a quantum computer cluster will break 256 bit keys before 2050** even if an adversary has either 2,8 billion 'commercial' quantum computers or 37 thousand 'worst case' quantum computers assuming:
   - Lenstra and Verheul's model for cryptanalytic progress;
   - no structural (quantum) attacks on the used cipher;
   - an overhead of 100x for quantum error correction in commercial quantum computers;
   - 1 operation per Grover iteration.

In conclusion, there are existing and widely used symmetric (AES-256) and hash primitives (SHA-2 and SHA-3 with an output length of at least 256 bits) that will withstand quantum computer attacks until way after 2050. For these primitives, it is not necessary to look for alternatives with the same urgency as the search for quantum-safe asymmetric primitives.

However, industry does have to worry about symmetric algorithms, since there are billions of devices in the world that rely on a symmetric cipher with a key length of 128 bits or less. The calculations used in the present document to assert that AES-256 will remain secure until way after 2050 cannot be used to predict when a quantum computer can attack AES-128, or any other cipher with a short key length. Therefore, industry is advised to identify where their products rely on smaller key and hash output lengths, and to start investigating the necessary steps for a transition to primitives with key lengths that will withstand quantum computer attacks like the ones investigated in the present document.

# History

| Document history | | |
|---|---|---|
| V1.1.1 | February 2017 | Publication |
| | | |
| | | |
| | | |
| | | |