

Internet Draft
Category: Informational
Document: [draft-mraihi-oath-hmac-otp-00.txt](#)
Expires: April 2005

D. MæRaihi
Verisign
M. Bellare
UCSD
F. Hoornaert
Vasco
D. Naccache
Gemplus
O. Ranen
Aladdin
October 2004

HOTP: An HMAC-based One Time Password Algorithm

Status of this Memo

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](#).

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>

Abstract

This document describes an algorithm to generate one-time password values, based on HMAC [[BCK1](#)]. A security analysis of the algorithm is presented, and important parameters related to the secure deployment of the algorithm are discussed. The proposed algorithm can be used across a wide range of network applications ranging from remote VPN access, Wi-Fi network logon to transaction-oriented Web applications.

This work is a joint effort by the OATH (Open AuTHentication) membership to specify an algorithm that can be freely distributed to the technical community. The authors believe that a common and

HOTP: An HMAC-based One Time Password Algorithm

October 2004

shared algorithm will facilitate adoption of two-factor authentication on the Internet by enabling interoperability across commercial and open-source implementations.

Table of Contents

1.	Overview.....	2
2.	Introduction.....	3
3.	Requirements Terminology.....	4
4.	Algorithm Requirements.....	4
5.	HOTP Algorithm.....	5
5.1	Notation and Symbols.....	5
5.2	Description.....	6
5.3	Generating an HOTP value.....	6
5.4	Example of HOTP computation for Digit = 6.....	7
6.	Security and Deployment Considerations.....	8
6.1	Authentication Protocol Requirements.....	8
6.2	Validation of HOTP values.....	8
6.3	Throttling at the server.....	9
6.4	Resynchronization of the counter.....	9
7.	HOTP Algorithm Security: Overview.....	10
8.	Protocol Extensions and Improvements.....	10
8.1	Number of Digits.....	10
8.2	Alpha-numeric Values.....	11
8.3	Sequence of HOTP values.....	11
8.4	A Counter-based Re-Synchronization Method.....	12
9.	Conclusion.....	12
10.	Acknowledgements.....	13
11.	Contributors.....	13
12.	References.....	13
12.1	Normative.....	13
12.2	Informative.....	13
13.	Authors' Addresses.....	13
	Appendix - HOTP Algorithm Security: Detailed Analysis.....	14
A.1	Definitions and Notations.....	14
A.2	The idealized algorithm: HOTP-IDEAL.....	15
A.3	Model of Security.....	15
A.4	Security of the ideal authentication algorithm.....	17
A.4.1	From bits to digits.....	17
A.4.2	Brute force attacks.....	18
A.4.3	Brute force attacks are the best possible attacks.....	19
A.5	Security Analysis of HOTP.....	20

1. Overview

The document introduces first the context around the HOTP algorithm. In [section 4](#), the algorithm requirements are listed and in [section 5](#), the HOTP algorithm is described. Sections [6](#) and [7](#) focus on the algorithm security. [Section 8](#) proposes some extensions

and improvements, and [Section 9](#) concludes this document. The interested reader will find in the Appendix a detailed, full-fledge analysis of the algorithm security: an idealized version of the algorithm is evaluated, and then the HOTP algorithm security is analyzed.

2. Introduction

Today, deployment of two-factor authentication remains extremely limited in scope and scale. Despite increasingly higher levels of threats and attacks, most Internet applications still rely on weak authentication schemes for policing user access. The lack of interoperability among hardware and software technology vendors has been a limiting factor in the adoption of two-factor authentication technology. In particular, the absence of open specifications has led to solutions where hardware and software components are tightly coupled through proprietary technology, resulting in high cost solutions, poor adoption and limited innovation.

In the last two years, the rapid rise of network threats has exposed the inadequacies of static passwords as the primary mean of authentication on the Internet. At the same time, the current approach that requires an end-user to carry an expensive, single-function device that is only used to authenticate to the network is clearly not the right answer. For two factor authentication to propagate on the Internet, it will have to be embedded in more flexible devices that can work across a wide range of applications.

The ability to embed this base technology while ensuring broad interoperability require that it be made freely available to the broad technical community of hardware and software developers. Only an open system approach will ensure that basic two-factor authentication primitives can be built into the next-generation of consumer devices such USB mass storage devices, IP phones, and personal digital assistants).

One Time Password is certainly one of the simplest and most popular forms of two-factor authentication for securing network access. For example, in large enterprises, Virtual Private Network access often requires the use of One Time Password tokens for remote user authentication. One Time Passwords are often preferred to stronger forms of authentication such as PKI or biometrics because an air-gap device does not require the installation of any client desktop software on the user machine, therefore allowing them to roam across multiple machines including home computers, kiosks and personal digital assistants.

This draft proposes a simple One Time Password algorithm that can be implemented by any hardware manufacturer or software developer

to create interoperable authentication devices and software agents. The algorithm is event-based so that it can be embedded in high volume devices such as Java smart cards, USB dongles and GSM SIM cards. The presented algorithm is made freely available to the developer community under the terms and conditions of the IETF Intellectual Property Rights [[RFC3668](#)].

The authors of this document are members of the Open AuTHentication initiative [[OATH](#)]. The initiative was created in 2004 to facilitate collaboration among strong authentication technology providers.

3. Requirements Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#).

4. Algorithm Requirements

This section presents the main requirements that drove this algorithm design. A lot of emphasis was placed on end-consumer usability as well as the ability for the algorithm to be implemented by low cost hardware that may provide minimal user interface capabilities. In particular, the ability to embed the algorithm into high volume SIM and Java cards was a fundamental pre-requisite.

R1 - The algorithm MUST be sequence or counter-based: One of the

goals is to have the HOTP algorithm embedded in high volume devices such as Java smart cards, USB dongles and GSM SIM cards.

R2 - The algorithm SHOULD be economical to implement in hardware by minimizing requirements on battery, number of buttons, computational horsepower, and size of LCD display. The algorithm MUST work with tokens that do not supports any numeric input, but MAY also be used with more sophisticated devices such as secure PIN-pads.

R3 - The value displayed on the token MUST be easily read and entered by the user: This requires the HOTP value to be of reasonable length. The HOTP value must be at least a 6-digit value. It is also desirable that the HOTP value be 'numeric only' so that it can be easily entered on restricted devices such as phones.

R4 - There MUST be user-friendly mechanisms available to resynchronize the counter. The sections [6.4](#) and [8.4](#) detail the resynchronization mechanism proposed in this draft.

R5- The algorithm MUST use a strong shared secret. The length of the shared secret MUST be at least 128 bits. This draft RECOMMENDs a shared secret length of 160 bits.

5. HOTP Algorithm

In this section, we introduce the notation and describe the HOTP algorithm basic blocks û the base function to compute an HMAC-SHA-1 value and the truncation method to extract an HOTP value.

5.1 Notation and Symbols

A string always means a binary string, meaning a sequence of zeros and ones.

If s is a string then $|s|$ denotes its length.

If n is a number then $|n|$ denotes its absolute value.

If s is a string then $s[i]$ denotes its i -th bit. We start numbering the bits at 0, so $s = s[0]s[1]..s[n-1]$ where $n = |s|$ is the length

of s .

Let StToNum (String to Number) denote the function which as input a string s returns the number whose binary representation is s . (For example $\text{StToNum}(110) = 6$).

Here is a list of symbols used in this document.

Symbol	Represents
--------	------------

C	8-byte (Little Endian) counter value, which is the modulo factor. This counter MUST be synchronized between the HOTP generator (client) and the HOTP validator (server);
---	--

K	shared secret between the client and the server; each HOTP generator has a different and unique secret K ;
---	--

T	throttling parameter: the server will refuse connections from a user after T unsuccessful authentication attempts;
---	--

s	resynchronization parameter: the server will attempt to verify a received authenticator across s consecutive counter values;
---	--

Digit	number of digits in an HOTP value; system parameter.
-------	--

5.2 Description

The HOTP algorithm is based on an increasing counter value and a static symmetric key known only to the token and the validation service. In order to create the HOTP value, we will use the HMAC-SHA-1 algorithm, as defined in [RFC 2104](#) [[BCK2](#)].

As the output of the HMAC-SHA1 calculation is 160 bits, we must truncate this value to something that can be easily entered by a user.

$$\text{HOTP}(K, C) = \text{Truncate}(\text{HMAC-SHA-1}(K, C))$$

Where:

- Truncate represents the function that converts an HMAC-SHA-1 value into an HOTP value as defined in [Section 5.3](#).

The HOTP values generated by the HOTP generator are treated as big endian.

5.3 Generating an HOTP value

We can describe the operations in 3 distinct steps:

Step 1: Generate an HMAC-SHA-1 value

```
Let HS = HMAC-SHA-1(K,C) // HS is a 20 byte string
```

Step 2: Generate a 4-byte string (Dynamic Truncation)

```
Let Sbits = DT(HS) // DT, defined in Section 6.3.1  
// returns a 31 bit string
```

Step 3: Compute an HOTP value

```
Let Snum = StToNum(S) // Convert S to a number in  
// 0...2{31}-1  
Return D = Snum mod 10Digit // D is a number in the range  
// 0...10{Digit}-1
```

The Truncate function performs Step 2 and Step 3, i.e. the dynamic truncation and then the reduction modulo 10^{Digit}. The purpose of the dynamic offset truncation technique is to extract a 4-byte dynamic binary code from a 160-bit (20-byte) HMAC-SHA1 result.

```
DT(String) // String = String[0]...String[19]  
Let OffsetBits be the low order four bits of String[19]  
Offset = StToNum(OffsetBits) // 0 <= Offset <= 15  
Let P = String[Offset]...String[Offset+3]  
Return the Last 31 bits of P
```

The reason for masking the most significant bit of P is to avoid confusion about signed vs. unsigned modulo computations. Different processors perform these operations differently, and masking out the signed bit removes all ambiguity.

Implementations MUST extract a 6-digit code at a minimum and possibly 7 and 8-digit code. Depending on security requirements,

Digit = 7 or more SHOULD be considered in order to extract a longer HOTP value.

The following paragraph is an example of using this technique for Digit = 6, i.e. that a 6-digit HOTP value is calculated from the HMAC value.

5.4 Example of HOTP computation for Digit = 6

The following code example describes the extraction of a dynamic binary code given that hmac_result is a byte array with the HMAC-SHA1 result:

```
int offset    = hmac_result[19] & 0xf ;
int bin_code  = (hmac_result[offset] & 0x7f) << 24
    | (hmac_result[offset+1] & 0xff) << 16
    | (hmac_result[offset+2] & 0xff) <<  8
    | (hmac_result[offset+3] & 0xff) ;
```

SHA-1 HMAC Bytes (Example)

Byte Number

00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19

Byte Value

1f 86 98 69 0e 02 ca 16 61 85 50 ef 7f 19 da 8e 94 5b 55 5a
-----*****-----++

- * The last byte (byte 19) has the hex value 0x5a.
- * The value of the lower four bits is 0xa (the offset value).
- * The offset value is byte 10 (0xa).
- * The value of the 4 bytes starting at byte 10 is 0x50ef7f19, which is the dynamic binary code DBC1
- * The MSB of DBC1 is 0x50 so DBC2 = DBC1 = 0x50ef7f19
- * HOTP = DBC2 modulo 10^6 = 872921.

We treat the dynamic binary code as a 31-bit, unsigned, big-endian integer; the first byte is masked with a 0x7f.

We then take this number modulo 1,000,000 (10^6) to generate the 6-digit HOTP value 872921 decimal.

6. Security and Deployment Considerations

Any One-Time Password algorithm is only as secure as the application and the authentication protocols that implement it. Therefore, this section discusses the critical security requirements that our choice of algorithm imposes on the authentication protocol and validation software. The parameters T and s discussed in this section have a significant impact on the security — further details in [Section 7](#) elaborate on the relations between these parameters and their impact on the system security.

6.1 Authentication Protocol Requirements

We introduce in this section some requirements for a protocol P implementing HOTP as the authentication method between a prover and a verifier.

RP1 – P MUST be two-factor, i.e. something you know (secret code such as a Password, Pass phrase, PIN code, etc.) and something you have (token). The secret code is known only to the user and usually entered with the one-time password value for authentication purpose (two-factor authentication).

RP3 – P MUST NOT be vulnerable to brute force attacks. This implies that a throttling/lockout scheme is REQUIRED on the validation server side.

RP4 – P SHOULD be implemented with respect to the state of the art in terms of security, in order to avoid the usual attacks and risks associated with the transmission of sensitive data over a public network (privacy, replay attacks, etc.)

6.2 Validation of HOTP values

The HOTP client (hardware or software token) increments its counter and then calculates the next HOTP value $HOTP_{client}$. If the value received by the authentication server matches the value calculated by the client, then the HOTP value is validated. In this case, the server increments the counter value by one.

If the value received by the server does not match the value calculated by the client, the server initiate the resynch protocol (look-ahead window) before it requests another pass.

If the resynch fails, the server asks then for another authentication pass of the protocol to take place, until the maximum number of authorized attempts is reached.

If and when the maximum number of authorized attempts is reached, the server SHOULD lock out the account and initiate a procedure to inform the user.

6.3 Throttling at the server

Truncating the HMAC-SHA1 value to a shorter value makes a brute force attack possible. Therefore, the authentication server needs to detect and stop brute force attacks.

We RECOMMEND setting a throttling parameter T , which defines the maximum number of possible attempts for One-Time-Password validation. The validation server manages individual counters per HOTP device in order to take note of any failed attempt. We RECOMMEND T not to be too large, particularly if the resynchronization method used on the server is window-based, and the window size is large. T SHOULD be set as low as possible, while still ensuring usability is not significantly impacted.

6.4 Resynchronization of the counter

Although the server's counter value is only incremented after a successful HOTP authentication, the counter on the token is incremented every time a new HOTP is requested by the user. Because of this, the counter values on the server and on the token might be out of synchronization.

We RECOMMEND setting a look-ahead parameter s on the server, which defines the size of the look-ahead window. In a nutshell, the server can recalculate the next s HOTP-server values, and check them against the received HOTP-client.

Synchronization of counters in this scenario simply requires the server to calculate the next HOTP values and determine if there is a match. Optionally, the system MAY require the user to send a sequence of (say 2, 3) HOTP values for resynchronization purpose, since forging a sequence of consecutive HOTP values is even more difficult than guessing a single HOTP value.

The upper bound set by the parameter s ensures the server does not go on checking HOTP values forever (causing a DoS attack) and also restricts the space of possible solutions for an attacker trying to

manufacture HOTP values. s SHOULD be set as low as possible, while still ensuring usability is not impacted.

7. HOTP Algorithm Security: Overview

The conclusion of the security analysis detailed in the Appendix section is that, for all practical purposes, the outputs of HOTP on distinct counter inputs are uniformly and independently distributed strings.

As a result, the best possible attack against the HOTP function is the brute force attack.

Assuming an adversary is able to observe numerous protocol exchanges and collect sequences of successful authentication values. This adversary, trying to build a function F to generate HOTP values based on his observations, will not have a significant advantage over a random guess.

The logical conclusion is simply that is best strategy will once again be to perform a brute force attack to enumerate and try all the possible values.

Considering the security analysis in the Appendix section of this document, without loss of generality, we can approximate closely the security of the HOTP algorithm by the following formula:

$$\text{Sec} = sv/10^{\text{Digit}}$$

Where:

- Sec is the probability of success of the adversary
- s stands for the look-ahead synchronization window size;
- v stands for the number of verification attempts;
- Digit stands for the number of digits in HOTP values.

Obviously, we can play with s , T (the Throttling parameter that would limit the number of attempts by an attacker) and Digit until achieving a certain level of security, still preserving the system usability.

[8. Protocol Extensions and Improvements](#)

We introduce in this section several enhancements and suggestions to further improve the security of the algorithm HOTP

8.1 Number of Digits

A simple enhancement in terms of security would be to extract more digits from the HMAC-SHA1 value.

For instance, calculating the HOTP value modulo 10^8 to build an 8-digit HOTP value would reduce the probability of success of the adversary from $sv/10^6$ to $sv/10^8$.

This could give the opportunity to improve usability, e.g. by increasing T and/or s , while still achieving a better security overall. For instance, $s = 10$ and $10v/10^8 = v/10^7 < v/10^6$ which is the theoretical optimum for 6-digit code when $s = 1$.

8.2 Alpha-numeric Values

Another option is to use A-Z and 0-9 values; or rather a subset of 32 symbols taken from the alphanumerical alphabet in order to avoid any confusion between characters: 0, O and Q as well as l, 1 and I are very similar, and can look the same on a small display.

The immediate consequence is that the security is now in the order of $sv/32^6$ for a 6-digit HOTP value and $sv/32^8$ for an 8-digit HOTP value.

$32^6 > 10^9$ so the security of a 6-alphanumeric HOTP code is slightly better than a 9-digit HOTP value, which is the maximum length of an HOTP code supported by the proposed algorithm.

$32^8 > 10^{12}$ so the security of an 8-alphanumeric HOTP code is significantly better than a 9-digit HOTP value.

Depending on the application and token/interface used for displaying and entering the HOTP value, the choice of alphanumeric values could be a simple and efficient way to improve security at a reduced cost and impact on users.

8.3 Sequence of HOTP values

As we suggested for the resynchronization to enter a short sequence (say 2 or 3) of HOTP values, we could generalize the concept to the protocol, and add a parameter *L* that would define the length of the HOTP sequence to enter.

Per default, the value *L* SHOULD be set to 1, but if security needs to be increased, users might be asked (possibly for a short period of time, or a specific operation) to enter *L* HOTP values.

This is another way, without increasing the HOTP length or using alphanumeric values to tighten security.

Note: The system MAY also be programmed to request synchronization on a regular basis (e.g. every night, or twice a week, etc.) and to achieve this purpose, ask for a sequence of *L* HOTP values.

8.4 A Counter-based Re-Synchronization Method

In this case, we assume that the client can access and send not only the HOTP value but also other information, more specifically the counter value.

A more efficient and secure method for resynchronization is possible in this case. The client application will not send the HOTP-client value only, but the HOTP-client and the related C-client counter value, the HOTP value acting as a message authentication code of the counter.

Resynchronization Counter-based Protocol (RCP)

The server accepts if the following are all true, where *C-server* is its own current counter value:

- 1) $C\text{-client} \geq C\text{-server}$
- 2) $C\text{-client} \hat{=} C\text{-server} \leq s$
- 3) Check that HOTP-client is valid $\text{HOTP}(K, C\text{-Client})$
- 4) If true, the server sets *C* to *C-client* + 1 and client is authenticated

In this case, there is no need for managing a look-ahead window anymore. The probability of success of the adversary is only $v/10^6$ or roughly v in one million. A side benefit is obviously to be able to increase s infinitely and therefore improve the system usability without impacting the security.

This resynchronization protocol SHOULD be use whenever the related impact on the client and server applications is deemed acceptable.

9. Conclusion

This draft describes HOTP, a HMAC-based One-Time Password algorithm. It also recommends the preferred implementation and related modes of operations for deploying the algorithm.

The draft also exhibits elements of security and demonstrates that the HOTP algorithm is practical and sound, the best possible attack being a brute force attack that can be prevented by careful implementation of countermeasures in the validation server.

Eventually, several enhancements have been proposed, in order to improve security if needed for specific applications.

10. Acknowledgements

The authors would like to thank Siddharth Bajaj, Alex Deacon and Nico Popp for their help during the conception and redaction of this document.

11. Contributors

The authors of this draft would like to emphasize the role of two persons who have made a key contribution to this document:

- Laszlo Elteto is system architect with SafeNet, Inc.
- Ernesto Frutos is director of Engineering with Authenex, Inc.

Without their advice and valuable inputs, this draft would not be the same.

12. References

12.1 Normative

- [BCK1] M. Bellare, R. Canetti, and H. Krawczyk, Keyed Hash Functions and Message Authentication, Proceedings of Crypto'96, LNCS Vol. 1109, pp. 1-15.
- [BCK2] M. Bellare, R. Canetti, and H. Krawczyk, HMAC: Keyed-Hashing for Message Authentication, IETF Network Working Group, [RFC 2104](#), February 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3668] Bradner, S., "Intellectual Property Rights in IETF Technology", [BCP 79](#), [RFC 3668](#), February 2004.

12.2 Informative

- [OATH] www.openauthentication.org, Initiative for Open AuThentication

13. Authors' Addresses

Primary point of contact (for sending comments and question):

David M'Raihi
VeriSign, Inc.
685 E. Middlefield Road Phone: 1-650-426-3832
Mountain View, CA 94043 USA Email: dmraihi@verisign.com

Other Authors' contact information:

Mihir Bellare
Dept of Computer Science and Engineering, Mail Code 0114
University of California at San Diego
9500 Gilman Drive
La Jolla, CA 92093, USA Email: mihir@cs.ucsd.edu

Frank Hoornaert
VASCO Data Security, Inc.

Koningin Astridlaan 164
1780 Wemmel, Belgium

Email: frh@vasco.com

David Naccache
Gemplus Innovation
34 rue Guynemer, 92447,
Issy les Moulineaux, France
and
Information Security Group,
Royal Holloway,
University of London, Egham,
Surrey TW20 0EX, UK

Email: david.naccache@gemplus.com

Email: david.naccache@rhul.ac.uk

Ohad Ranen
Aladdin Knowledge Systems Ltd.
15 Beit Oved Street
Tel Aviv, Israel 61110

Email: Ohad.Ranen@ealaddin.com

Appendix – HOTP Algorithm Security: Detailed Analysis

The security analysis of the HOTP algorithm is summarized in this section. We first detail the best attack strategies, and then elaborate on the security under various assumptions, the impact of the truncation and some recommendations regarding the number of digits.

We focus this analysis on the case where $\text{Digit} = 6$, i.e. an HOTP function that produces 6-digit values, which is the bare minimum recommended in this draft.

A.1 Definitions and Notations

We denote by $\{0,1\}^l$ the set of all strings of length l .

Let $Z_{\{n\}} = \{0, \dots, n - 1\}$.

Let $\text{IntDiv}(a,b)$ denote the integer division algorithm that takes input integers a, b where $a \geq b \geq 1$ and returns integers (q,r) the quotient and remainder, respectively, of the division of a by

b. (Thus $a = bq + r$ and $0 \leq r < b$.)

Let $H: \{0,1\}^k \times \{0,1\}^c \rightarrow \{0,1\}^n$ be the base function that takes a k -bit key K and c -bit counter C and returns an n -bit output $H(K,C)$. (In the case of HOTP, H is HMAC-SHA-1; we use this formal definition for generalizing our proof of security)

A.2 The idealized algorithm: HOTP-IDEAL

We now define an idealized counterpart of the HOTP algorithm. In this algorithm, the role of H is played by a random function that forms the key.

To be more precise, let $\text{Maps}(c,n)$ denote the set of all functions mapping from $\{0,1\}^c$ to $\{0,1\}^n$. The idealized algorithm has key space $\text{Maps}(c,n)$, so that a "key" for such an algorithm is a function h from $\{0,1\}^c$ to $\{0,1\}^n$. We imagine this key (function) to be drawn at random. It is not feasible to implement this idealized algorithm, since the key, being a function from is way too large to even store. So why consider it?

Our security analysis will show that as long as H satisfies a certain well-accepted assumption, the security of the actual and idealized algorithms is for all practical purposes the same. The task that really faces us, then, is to assess the security of the idealized algorithm.

In analyzing the idealized algorithm, we are concentrating on assessing the quality of the design of the algorithm itself, independently of HMAC-SHA-1. This is in fact the important issue.

A.3 Model of Security

The model exhibits the type of threats or attacks that are being considered and enables to assess the security of HOTP and HOTP-IDEAL. We denote ALG as either HOTP or HOTP-IDEAL for the purpose of this security analysis.

The scenario we are considering is that a user and server share a key K for ALG . Both maintain a counter C , initially zero, and the user authenticates itself by sending $\text{ALG}(K,C)$ to the server. The latter accepts if this value is correct.

In order to protect against accidental increment of the user counter, the server, upon receiving a value z , will accept as long

as z equals $\text{ALG}(K,i)$ for some i in the range $C, \dots, C + s - 1$, where s is the resynchronization parameter and C is the server counter. If it accepts with some value of i , it then increments its counter to $i + 1$. If it does not accept, it does not change its counter value.

The model we specify captures what an adversary can do and what it needs to achieve in order to win. First, the adversary is assumed to be able to eavesdrop, meaning see the authenticator transmitted by the user. Second, the adversary wins if it can get the server to accept an authenticator relative to a counter value for which the user has never transmitted an authenticator.

The formal adversary, which we denote by B , starts out knowing which algorithm ALG is being used, knowing the system design and knowing all system parameters. The one and only thing it is not given a priori is the key K shared between the user and the server.

The model gives B full control of the scheduling of events. It has access to an authenticator oracle representing the user. By calling this oracle, the adversary can ask the user to authenticate itself and get back the authenticator in return. It can call this oracle as often as it wants and when it wants, using the authenticators it accumulates to perhaps learn how to make authenticators itself. At any time, it may also call a verification oracle, supplying the latter with a candidate authenticator of its choice. It wins if the server accepts this accumulator.

Consider the following game involving an adversary B that is attempting to compromise the security of an authentication algorithm $\text{ALG}: K \times \{0,1\}^c \rightarrow R$.

Initializations – A key K is selected at random from K , a counter C is initialized to 0, and the Boolean value `win` is set to false.

Game execution – Adversary B is provided with the two following oracles:

```
Oracle Auth0()
  O = ALG(K,C)
  C = C + 1
  Return O to B
```

```
Oracle Ver0()
  i = C
  While (i <= C + s - 1 and Win = FALSE) do
    If O = ALG(K,i) then Win = TRUE; C = i + 1
    Else i = i + 1
```

Auth0() is the authenticator oracle and Ver0() is the verification oracle.

Upon execution, B queries the two oracles at will. Let Adv(B) be the probability that win gets set to true in the above game. This is the probability that the adversary successfully impersonates the user.

Our goal is to assess how large this value can be as a function of the number v of verification queries made by B, the number a of authenticator oracle queries made by B, and the running time t of B. This will tell us how to set the throttle, which effectively upper bounds v .

A.4 Security of the ideal authentication algorithm

This section summarizes the security analysis of HOTP-IDEAL, starting with the impact of the conversion modulo 10^{Digit} and then, focusing on the different possible attacks.

A.4.1 From bits to digits

The dynamic offset truncation of a random n -bit string yields a random 31-bit string. What happens to the distribution when it is taken modulo $m = 10^{\text{Digit}}$, as done in HOTP?

The following lemma estimates the biases in the outputs in this case.

Lemma 1

Let $N \geq m \geq 1$ be integers, and let $(q,r) = \text{IntDiv}(N,m)$. For z in $Z_{\{m\}}$ let:

$$P_{\{N,m\}}(z) = \Pr [x \bmod m = z : x \text{ randomly pick in } Z_{\{n\}}]$$

Then for any z in $Z_{\{m\}}$

$$P_{\{N,m\}}(z) = \begin{array}{ll} (q + 1) / N & \text{if } 0 \leq z < r \\ q / N & \text{if } r \leq z < m \end{array}$$

Proof of Lemma 1

Let the random variable X be uniformly distributed over $Z_{\{N\}}$.
Then:

$$\begin{aligned}
 P_{\{N,m\}}(z) &= \Pr [X \bmod m = z] \\
 &= \Pr [X < mq] \cdot \Pr [X \bmod m = z \mid X < mq] \\
 &\quad + \Pr [mq \leq X < N] \cdot \Pr [X \bmod m = z \mid mq \leq X < N]
 \end{aligned}$$

$$\begin{aligned}
 &= \frac{mq}{N} \cdot \frac{1}{m} + \frac{(N - mq)}{N} \cdot \frac{1}{N - mq} && \begin{aligned} &\text{if } 0 \leq z < N - mq \\ &\text{if } N - mq \leq z \leq N \end{aligned} \\
 &= \frac{q}{N} + \frac{r}{N} \cdot \frac{1}{r} && \begin{aligned} &\text{if } 0 \leq z < N - mq \\ &\text{if } r \leq z \leq N \end{aligned}
 \end{aligned}$$

Simplifying yields the claimed equation.

Let $N = 2^{31}$, $d = 6$ and $m = 10^d$. If x is chosen at random from $Z_{\{N\}}$ (meaning, is a random 31-bit string), then reducing it to a 6-digit number by taking $x \bmod m$ does not yield a random 6-digit number.

Rather, $x \bmod m$ is distributed as shown in the following table:

Values	Probability that each appears as output
0,1,...,483647	$2148/2^{31}$ roughly equals to $1.00024045/10^6$
483648,...,999999	$2147/2^{31}$ roughly equals to $0.99977478/10^6$

If X is uniformly distributed over $Z_{\{2^{31}\}}$ (meaning is a random 31-bit string) then the above shows the probabilities for different outputs of $X \bmod 10^6$. The first set of values appear with probability slightly greater than 10^{-6} , the rest with probability slightly less, meaning the distribution is slightly non-uniform.

However, as the Figure indicates, the bias is small and as we will see later, negligible: the probabilities are very close to 10^{-6} .

A.4.2 Brute force attacks

If the authenticator consisted of d random digits, then a brute force attack using v verification attempts would succeed with probability $sv/10^{\text{Digit}}$.

However, an adversary can exploit the bias in the outputs of HOTP-IDEAL, predicted by Lemma 1, to mount a slightly better attack.

Namely, it makes authentication attempts with authenticators which are the most likely values, meaning the ones in the range $0, \dots, r - 1$, where $(q, r) = \text{IntDiv}(2^{31}, 10^{\text{Digit}})$.

The following specifies an adversary in our model of security that mounts the attack. It estimates the success probability as a function of the number of verification queries.

For simplicity, we assume the number of verification queries is at most r . With $N = 2^{31}$ and $m = 10^6$ we have $r = 483,648$, and the throttle value is certainly less than this, so this assumption is not much of a restriction.

Proposition 1

Suppose $m = 10^{\text{Digit}} < 2^{31}$, and let $(q, r) = \text{IntDiv}(2^{31}, m)$. Assume $s \leq m$. The brute-force attack adversary $B\text{-bf}$ attacks HOTP using $v \leq r$ verification oracle queries. This adversary makes no authenticator oracle queries, and succeeds with probability

$$\text{Adv}(B\text{-bf}) = 1 - (1 - v(q+1)/2^{31})^s$$

which is roughly equals to

$$sv, (q+1)/2^{31}$$

With $m = 10^6$ we get $q = 2,147$. In that case, the brute force attack using v verification attempts succeeds with probability

$$\text{Adv}(B\text{-bf}) \text{ roughly } = sv, 2148/2^{31} = sv, 1.00024045/10^6$$

As this equation shows, the resynchronization parameter s has a significant impact in that the adversary's success probability is proportional to s . This means that s cannot be made too large without compromising security.

A.4.3 Brute force attacks are the best possible attacks

A central question is whether there are attacks any better than the brute force one. In particular, the brute force attack did not attempt to collect authenticators sent by the user and try to cryptanalyze them in an attempt to learn how to better construct authenticators. Would doing this help? Is there some way to learn how to build authenticators that result in a higher success rate than given by the brute-force attack?

The following says the answer to these questions is no. No matter what strategy the adversary uses, and even if it sees, and tries to exploit, the authenticators from authentication attempts of the user, its success probability will not be above that of the brute force attack – this is true as long as the number of authentications it observes is not incredibly large. This is valuable information regarding the security of the scheme.

Proposition 2

Suppose $m = 10^{\text{Digit}} < 2^{31}$, and let $(q,r) = \text{IntDiv}(2^{31},m)$. Let B be any adversary attacking HOTP-IDEAL using v verification oracle queries and $a \leq 2^c$ authenticator oracle queries. Then

$$\text{Adv}(B) \leq \frac{sv}{2^{31}}, (q+1)$$

Note: This result is conditional on the adversary not seeing more than 2^c authentications performed by the user, which is hardly restrictive as long as c is large enough.

With $m = 10^6$ we get $q = 2,147$. In that case, Proposition 2 says that any adversary B attacking HOTP-IDEAL and making v verification attempts succeeds with probability at most

Equation 1

$$\frac{sv}{2^{31}}, 2148 \text{ roughly} = \frac{sv}{1.00024045/10^6}$$

Meaning, B 's success rate is not more than that achieved by the brute force attack.

A.5 Security Analysis of HOTP

We have analyzed in the previous sections, the security of the idealized counterparts HOTP-IDEAL of the actual authentication algorithm HOTP. We now show that, under appropriate and well-believed assumption on H , the security of the actual algorithms is essentially the same as that of its idealized counterpart.

The assumption in question is that H is a secure pseudorandom function, or PRF, meaning that its input-output values are indistinguishable from those of a random function in practice.

Consider an adversary A that is given an oracle for a function $f: \{0,1\}^c \rightarrow \{0,1\}^n$ and eventually outputs a bit. We denote $\text{Adv}(A)$ as the prf-advantage of A , which represents how well the adversary does at distinguishing the case where its oracle is $H(K, \cdot)$ from the case where its oracle is a random function of $\{0,1\}^c$ to $\{0,1\}^n$.

One possible attack is based on exhaustive search for the key K . If A runs for t steps and T denotes the time to perform one computation of H , its prf-advantage from this attack turns out to be $(t/T)2^{-k}$. Another possible attack is a birthday one [3], whereby A can attain advantage $p^2/2^n$ in p oracle queries and running time about pT .

Our assumption is that these are the best possible attacks. This translates into the following.

Assumption 1

Let T denotes the time to perform one computation of H . Then if A is any adversary with running time at most t and making at most p oracle queries,

$$\text{Adv}(A) \leq (t/T)/2^k + p^2/2^n$$

In practice this assumption means that H is very secure as PRF. For example, given that $k = n = 160$, an attacker with running time 2^{60} and making 2^{40} oracle queries has advantage at most (about) 2^{-80} .

Theorem 1

Suppose $m = 10^{\text{Digit}} < 2^{31}$, and let $(q,r) = \text{IntDiv}(2^{31},m)$. Let B be any adversary attacking HOTP using v verification oracle queries, $a \leq 2^c - s$ authenticator oracle queries, and running time t . Let T denote the time to perform one computation of H . If

Assumption 1 is true then

$$\text{Adv}(B) \leq sv \cdot (q + 1)/2^{31} + (t/T)/2^k + ((sv + a)^2)/2^n$$

In practice, the $(t/T)2^{-k} + ((sv + a)^2)2^{-n}$ term is much smaller than the $sv(q + 1)/2^n$ term, so that the above says that for all practical purposes the success rate of an adversary attacking HOTP is $sv(q + 1)/2^n$, just as for HOTP-IDEAL, meaning the HOTP algorithm is in practice essentially as good as its idealized counterpart.

In the case $m = 10^6$ of a 6-digit output this means that an adversary making v authentication attempts will have a success rate that is at most that of Equation 1.

For example, consider an adversary with running time at most 2^{60} that sees at most 2^{40} authentication attempts of the user. Both these choices are very generous to the adversary, who will typically not have these resources, but we are saying that even such a powerful adversary will not have more success than indicated by Equation 1.

We can safely assume $sv \leq 2^{40}$ due to the throttling and bounds on s . So:

$$(t/T)/2^k + ((sv + a)^2)/2^n \leq 2^{60}/2^{160} + (2^{41})^2/2^{160} \\ \text{roughly } \leq 2^{-78}$$

which is much smaller than the success probability of Equation 1 and negligible compared to it.

Full Copyright Statement

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.