Mathematical Mesh Part VI: The Trust Mesh
draft-hallambaker-mesh-trust-02

Abstract

   This paper extends Shannon's concept of a 'work factor' as applied to
   evaluation of cryptographic algorithms to provide an objective
   measure of the practical security offered by a protocol or
   infrastructure design. Considering the hypothetical work factor
   based on an informed estimate of the probable capabilities of an
   attacker with unknown resources provides a better indication of the
   relative strength of protocol designs than the computational work
   factor of the best-known attack.

   The social work factor is a measure of the trustworthiness of a
   credential issued in a PKI based on the cost of having obtained the
   credential through fraud at a certain point in time. Use of the
   social work factor allows evaluation of Certificate Authority based
   trust models and peer to peer (Web of Trust) models to be evaluated
   in the same framework. The analysis demonstrates that both
   approaches have limitations and that in certain applications, a
   blended model is superior to either by itself.

   The final section of the paper describes a proposal to realize this
   blended model using the Mathematical Mesh.

   This document is also available online at
   http://mathmesh.com/Documents/draft-hallambaker-mesh-trust.html [1] .

Status of This Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at https://datatracker.ietf.org/drafts/current/.

   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 as "work in progress."

This Internet-Draft will expire on January 9, 2020.

Copyright Notice

Table of Contents

1. Work Factor

Recent events have highlighted both the need for open standards-based
security protocols and the possibility that the design of such
protocols may have been sabotaged [Schneier2013] . We thus face two
important and difficult challenges, first to design an Internet
security infrastructure that offers practical security against the
class of attacks revealed, and secondly, to convince potential users
that the proposed new infrastructure has not been similarly
sabotaged.

The measure of a security of a system is the cost and difficulty of
making a successful attack. The security of a safe is measured by
the length time it is expected to resist attack using a specified set
of techniques. The security of a cryptographic algorithm against a
known attack is measured by the computational cost of the attack.

This paper extends Shannon's concept of a 'work factor' [Shannon1949]
to provide an objective measure of the security a protocol or
infrastructure offers against other forms of attack.

1.1. Computational Work Factor

The term 'Computational Work Factor' is used to refer to Shannon's
original concept.

One of Shannon's key insights was that the work factor of a
cryptographic algorithm could be exponential. Adding a single bit to
the key size of an ideal symmetric algorithm presents only a modest
increase in computational effort for the defender but doubles the
work factor for the attacker.

More precisely, the difficulty of breaking a cryptographic algorithm
is generally measured by the work-factor ratio. If the cost of
encrypting a block with 56-bit DES is x, the worst case cost of
recovering the key through a brute force attack is $2^{56}x$. The
security of DES has changed over time because x has fallen
exponentially.

While the work factor is traditionally measured in terms of the
number of operations, many cryptanalytic techniques permit memory
used to be traded for computational complexity. An attack requiring
2^{64} bytes of memory that reduces the number of operations required
to break a 128 bit cipher to 2^{64} is a rather lower concern than one
which reduces the number of operations to 2^{80}. The term 'cost' is
used to gloss over such distinctions.

The Computational Work Factor ratio WF-C (A) of a cryptographic
algorithm A, is the cost of the best-known attack divided by the cost
of the algorithm itself.

1.2. Hypothetical Work Factor

Modern cryptographic algorithms use keys of 128 bits or more and
present a work factor ratio of 2^{128} against brute force attack.
This work factor is at least 2^{72} times higher than DES and
comfortably higher than the work factor of 2^{80} operations that is
generally believed to be the practical limit to current attacks.

Though Moore's law has delivered exponential improvements in
computing performance over the past four decades, this has been
achieved through continual reductions in the minimum feature size of
VLSI circuits. As the minimum feature size rapidly approaches the
size of individual atoms, this mechanism has already begun to stall
[Intel2018] .

While an exceptionally well-resourced attacker may gain performance
advances through use of massive parallelism, faster clock rates made
possible by operating at super-low temperatures and custom designed
circuits, the return on such approaches is incremental rather than
exponential.

Performance improvements may allow an attacker to break systems with
a work factor several orders of magnitude greater than the public
state of the art. But an advance in cryptanalysis might permit a
potentially more significant reduction in the work factor.

The primary consideration in the choice of a cryptographic algorithm
therefore is not the known computational work factor as measured
according to the best publicly known attack but the confidence that
the computational work factor of the best attack that might be known
to the attacker.

While the exact capabilities of the adversary are unknown, a group of
informed experts may arrive at a conservative estimate of their
likely capabilities. In particular, it is the capabilities of
nation-state actors that generally give rise to greatest concern in

security protocol design. In this paper we refer to this set of
actors as nation-state class adversaries in recognition of the fact
that certain technology companies posses computing capabilities that
rival if not exceed those of the largest state actors and those
capabilities could at least in theory be co-opted for other purposes
in certain circumstances.

The probability that a nation-state class has discovered an attack
against AES-128 with a work factor ratio of 2^{120} might be considered
relatively high while the probability that an attack with a work
factor ratio of less than 2^{64} is very low.

We define the hypothetical work factor function WF-H (A, p) as
follows: If WF is a work factor ratio and p is an informed estimate
of the probability that an adversary has developed an attack with a
work factor ratio against algorithm A of WF or less then WF-H (A, p)
= WF.

Since the best-known public attack is known to the attacker, WF-H (A,
1) = _CWF (A)

The inverse function WF-H' (A, WF) returns the estimated probability
that the work factor of algorithm A is at least WF.

The hypothetical work factor and its inverse may be used to compare
the relative strengths of protocol designs. Given designs A and B,
we can state that B is an improvement on A if WF-H (A,p) > WF-H (B,p)
for all p.

When considering a protocol or infrastructure design we can thus
improve a protocol by either:

o Increasing WF-H (A,p) for some p, or

o Decreasing WF-H '(A,WF)

1.3. Known Unknowns

Unlike the computational work factor, the hypothetical work factor
does not provide an objective measure of the security offered by a
design. The purpose of the hypothetical work factor is to allow the
protocol designer to compare the security offered by different design
choices.

The task that the security engineer faces is to secure the system
from all attacks whether the attacks themselves are known or unknown.
In the current case it is known that an attacker is capable of

breaking at least some of the cryptographic algorithms in use. But
not which algorithms are affected or the nature of the attack(s).

Unlike the computational work factor, the hypothetical work factor
does not deliver an academically rigorous, publication and citation
worthy measure of the strength of a design. That is not its purpose.
the purpose of the hypothetical work factor is to assist the protocol
designer in designing protocols.

Design of security protocols has always required the designer to
consider attackers whose capabilities are not currently known and
thus involved a considerable degree of informed opinion and
guesswork. Whether correctly or not, the decision to reject changes
to the DNSSEC protocol to enable deployment in 2002 rested in part on
a statement by a Security Area Director that a proposed change gave
him 'a bad feeling in his gut'. The hypothetical work factor permits
the security designer to model to quantify such intestinally based
assumptions and model the effect on the security of the resulting
design.

Security is a property of systems rather than individual components.
While it is quite possible that there are no royal roads to
cryptanalysis and cryptanalysis of algorithms such as AES 128 is
infeasible even for the nation state class adversaries, such
adversaries are not limited to use of cryptanalytic attacks.

Despite the rise of organized cyber-crime, many financial systems
still employ weak cryptographic systems that are known to be
vulnerable to cryptanalytic attacks that are well within the
capabilities of the attackers. But fraud based on such techniques
remains vanishingly rare as it is much easier for the attackers to
persuade bank customers to simply give their access credentials to
the attacker.

Even if a nation-state class attacker has a factoring attack which
renders an attack on RSA-2048 feasible, it is almost certainly easier
for a nation-state class attacker to compromise a system using
RSA-2048 in other ways. For example, persuading the target of the
surveillance to use cryptographic devices with a random number
generator that leaks a crib for the attacker. Analyzing the second
form of attack requires a different type of analysis which is
addressed in the following section on social work factor.

1.4. Defense in Depth

The motivation behind introducing the concept of the hypothetical
work factor is a long experience of seeing attempts to make security

protocols more robust being deflected by recourse to specious
arguments based on the computational work factor.

For example, consider the case in which a choice between a single
security control and a defense in depth strategy is being considered:

o Option A: Uses algorithm X for protection.

o Option B: Uses a combination of algorithm X and algorithm Y for
   protection such that the attacker must defeat both to break the
   system and algorithms based on different cryptographic principles
   are chosen so as to minimize the risk of a common failure mode.

If the computational work factor for both algorithms X and Y is
2^{128}, both options present the same work factor ratio. Although
Option B offers twice the security, it also requires twice the work.

The argument that normally wins is that both options present the same
computational work factor ratio of 2^{128}, Option A is simpler and
therefore Option A should be chosen. This despite the obvious fact
that only Option B offers defense in depth.

If we consider the adversary of being capable of performing a work
factor ratio of 2^{80} and the probability the attacker has discovered
an attack capable of breaking algorithms X and Y to be 10% in each
case, the probability that the attacker can break Option A is 10%
while the probability that an attack on Option B is only 1%, a
significant improvement.

While Option B clearly offers a significant potential improvement in
security, this improvement is only fully realized if the
probabilities of a feasible attack are independent.

1.5. Mutual Reinforcement

The defense in depth approach affords a significant improvement in
security but an improvement that is incremental rather than
exponential in character. With mutual reinforcement we design the
mechanism such that in addition to requiring the attacker to break
each of the component algorithms, the difficulty of the attacks is
increased.

For example, consider the use of a Deterministic Random Number
Generator R(s,n) which returns a sequence of values R(s,1), R(s,2)...
from an initial seed s.

Two major concerns in the design of such generators are the possibility of bias and that the seed value be somehow leaked through a side channel.

Both concerns are mitigated if instead of using the output of one generator directly, two independent random number generators with distinct seeds are used.

For example, consider the use of the value R1(s1,n) XOR R2(s2,n) where R1(s,n) and R2(s,n) are different random number generation functions and s1, s2 are distinct seeds.

The XOR function has the property of preserving randomness so that the output is guaranteed to be at least as random as either of the generators from which it is built (provided that there is not a common failure mode). Further, recovery of either random seed is at least as hard as using the corresponding generator on its own. Thus, the Hypothetical work factor for the combined system is improved to at least the same extent as in the defense in depth case.

But any attempt to break either generator must now face the additional complexity introduced by the output being masked with the unknown output of the other. An attacker cannot cryptanalyze the two generator functions independently. If the two generators and the seeds are genuinely independent, the combined hypothetical work factor is the product of the hypothetical work factors from which it is built.

While implementing two independent generators and seeds represents a significant increase in cost for the implementer, a similar exponential leverage might be realized with negligible additional complexity through use of a cryptographic digest of the generator output to produce the masking value.

1.6. Safety in Numbers

In a traditional security analysis, the question of concern is whether a cryptanalytic attack is feasible or not. When considering an indiscriminate intercept capability as in a nation-state class attack, the concern is not just whether an individual communication might be compromised but the number of communications that may be compromised for a given amount of effort.

'Perfect' Forward Secrecy is an optional feature supported in IPSec and TLS. In 2008, implementations of TLS/1.2 [RFC6246] purported to offer a choice between:

Direct key exchange with a work factor dependent on the difficulty of breaking RSA 2048

Direct key exchange followed by a perfect forward secrecy exchange with a work factor dependent on the difficulty of breaking both RSA 2048 and DH 1024.

Using the computational work factor alone suggests that the second scheme has little advantage over the first since the computational work factor of Diffie Hellman using the best-known techniques 2^80 while the computational work factor for RSA 2048 is 2^112. Use of the perfect forward secrecy exchange has a significant impact on server performance but does not increase the difficulty of cryptanalysis.

Use of perfect forward secrecy with a combination of RSA and Diffie Hellman does not provide a significant improvement in the hypothetical work factor either if individual messages are considered. The RSA and Diffie Hellman systems are closely related and so an attacker that can break RSA 2048 can almost certainly break RSA 1024. Moreover, computational work factor for DH 1024 is only 2^80 and thus feasibly within the reach of a well-funded and determined attacker.

According to the analysis informally applied during design, use of perfect forward secrecy does provide an important security benefit when multiple messages are considered. While a sufficiently funded and determined attacker could conceivably break tens, hundreds or even thousands of DH 1024 keys a year, it is rather less likely that an attacker could break millions a year. The OCSP servers operated by Comodo CA receive over 2 billion hits a day and this represents only a fraction of the number of uses of TLS on the Internet. Use of perfect forward secrecy does not prevent an attacker from decrypting any particular message but raises the cost of indiscriminate intercept and decryption.

Unfortunately, this analysis is wrong because the TLS key exchange does not achieve a work factor dependent on the difficulty of breaking both RSA 2048 and DH 1024. The pre-master secret established in the initial RSA 2048 exchange is only used to authenticate the key exchange process itself. The session keys used to encrypt content are derived from the weaker ephemeral key exchange, the parameters of which are exchanged in plaintext. Due to this defect in the design of the protocol, the Work Factor of the protocol is the work factor of DH1024 alone.

Nor does the use of Diffie Hellman in this fashion provide security when multiple messages are exchanged. The Logjam attack [Adrian2015]

exploits the fact that the difficulty of breaking the discrete
logarithm involves four major steps, the first three of which are the
most computationally intensive and only depend on the shared group
parameters. The cost of breaking a hundred Diffie Hellman public
keys is not a hundred times the cost of breaking a single key, there
is almost no difference.

Work factor analysis exposes these flaws in the design of the
TLS/1.2. Since the session keys used to encrypt traffic do not
depend on knowing the secret established in the RSA2048 exchange, the
work factor of the protocol is the lesser of 2^80 and 2^112.

A simple means of ensuring that the work factor of a protocol is not
reduced by a fresh key exchange is to use a one-way function such as
a cryptographic digest or a key exchange to combine the output of the
prior exchange with its successor. This principle is employed in the
double ratchet algorithm [Ratchet] used in the Signal protocol. In
the Mesh, the HKDF Key Derivation function [RFC5869] is frequently
used for the same purpose.

The work factor downgrade issue was addressed in TLS/1.3 [RFC8446]
albeit in a less direct fashion by encrypting the ephemeral key
exchange.

1.7. Cost Factor

As previously discussed, cryptanalysis is not the only tool available
to an attacker. Faced with a robust cryptographic defense, Internet
criminals have employed 'social engineering' instead. A nation-state
class attacker may use any and every tool at their disposal including
tools that are unique to government backed adversaries such as the
threat of legal sanctions against trusted intermediaries.

Although attackers can and will use every tool at their disposal,
each tool carries a cost and some tools require considerable advance
planning to use. It is conceivable that the AES standard published
by NIST contains a backdoor that somehow escaped the extensive peer
review. But any such effort would have had to have begun well in
advance of 1998 when the Rijndael cipher was first published.

Nation-state class actors frequently rely for security on the same
infrastructures that they are attempting to attack. Thus, the
introduction of vulnerabilities that might also be exploited by the
opposition incurs a cost to both. This concern is recognized in the
NSA 'NOBUS' doctrine: Nobody but us. To introduce a vulnerability in
a random number generator that can only be exploited by a party that
knows the necessary private key is acceptable. But introducing a
vulnerability that depends on the use of an unpublished cryptanalytic

technique is not because that same technique might be discovered by
the opposition.

Subversion of cryptographic apparatus such as Hardware Security
Modules (HSMs) and SSL accelerators faces similar constraints. HSMs
may be compromised by an adversary but the compromise must have taken
place before the device was manufactured or serviced.

Just as computational attacks are limited by the cryptanalytic
techniques known to and the computational resources available to the
attacker, social attacks are limited by the cost of the attack and
the capacity of the attacker.

The Cost Factor $C(t)$ is an estimate of the cost of performing an
attack on or before a particular date in time (t).

For the sake of simplicity, currency units are used under the
assumption that all the resources required are fungible and that all
attackers face the same costs. But such assumptions may need to be
reconsidered when there is a range of attackers with very different
costs and capabilities. A hacktivist group could not conceivably
amass the computational and covert technical resources available to
the NSA but such a group could in certain circumstances conceivably
organize a protest with a million or more participants while the
number of NSA employees is believed to still be somewhat fewer.

The computational and hypothetical work factors are compared against
estimates of the computational resources of the attacker. An attack
is considered to be infeasible if that available computational
resources do not allow the attack to be performed within a useful
period of time.

The cost factor is likewise compared against an incentive estimate,
$I(t)$ which is also time based.

o   An attack is considered to be productive for an attacker if there
    was a time t for which $I(t) > C(t)$.

o   An attack is considered to be unproductive if there is no time at
    which it was productive for that attacker.

Unlike Cost Factor for which a lower bound based on the lowest cost
and highest capacity may be usefully applied to all attackers,
differences in the incentive estimate between attackers are likely to
be very significant. Almost every government has the means to
perform financial fraud on a vast scale but only rarely does a
government have the incentive. When governments do engage in

activities such as counterfeiting banknotes this has been done for
motives beyond mere peculation.

While government actors do not respond to the same incentives as
Internet criminals, governments fund espionage activities in the
expectation of a return on their investment. A government agency
director who does not produce the desired returns is likely to be
replaced.

For example, when the viability of SSL and the Web PKI for protecting
Internet payments was considered in the mid-1990s, the key question
was whether the full cost of obtaining a fraudulently issued
certificate would exceed the expected financial return where the full
cost is understood to include the cost of registering a bogus
corporation, submitting the documents and all the other activities
that would be required if a sustainable model for payments fraud was
to be established.

For an attack to be attractive to an attacker it is not just
necessary for it to be productive, the time between the initial
investment and the reward and the likelihood of success are also
important factors. An attack that requires several years of advance
planning is much less attractive than an attack which returns an
immediate profit.

An attack may be made less attractive by

o Increasing the cost

o Reducing the incentive

o Reducing the expected gain

o Reducing the probability that the incentive will be realized

o Increasing the time between the initial investment and the return.

Most real-world security infrastructures are based on more than one
of these approaches. The WebPKI is designed to increase the cost of
attack by introducing validation requirements and reduce the expected
gain through its revocation infrastructure.

1.8. Social Work Factor

In the cost factor analysis, it is assumed that all costs are
fungible, and the attack capacity of the attacker is only limited by
their financial resources. Some costs are not fungible however, in
particular inducing a large number of people to accept a forgery

without the effort being noticed requires much more than a limitless
supply of funds.

In a computational attack an operation will at worst fail to deliver
success. There is no penalty for failure beyond having failed to
succeed. When attempting to perpetuate a fraud on the general
public, every attempt carries a risk of exposure of the entire
scheme. When attempting to perform any covert activity, every
additional person who is indoctrinated into the conspiracy increases
the chance of exposure.

The totalitarian state envisioned by George Orwell in 1984 was only
plausible because each and every citizen is coerced to act as a party
to the conspiracy. The erasure and replacement of the past was
possible because the risk of exposure was nil.

In 2011, I expressed concern to a retired senior member of the NSA
staff that the number of contractors being hired to perform cyber-
sabotage operations represented a security risk and might be creating
a powerful constituency with an interest in the aggressive
militarization of cyberspace rather than preparing for its defense.
Subsequent disclosures by Robert Snowden have validated the
disclosure risk aspect of these concerns. Empirically, the NSA, an
organization charged with protecting the secrecy of government
documents, was unable to maintain the secrecy of their most important
secrets when the size of the conspiracy reached a few ten thousand
people.

The community of commercial practitioners of cryptographic
information security is small in size but encompasses many
nationalities. Many members of the community are bound by
ideological commitments to protecting personal privacy as an
unqualified moral objective.

Introducing a backdoor into a HSM, application or operating system
platform requires that every person with access to the platform
source or who might be called in to audit the code be a party to the
conspiracy. Tapping the fiber optic cables that support the Internet
backbone requires only a small work crew and digging equipment.
Maintaining a covert backdoor in a major operating system platform
would require hundreds if not thousands of engineers to participate
in the conspiracy.

The Social Work Factor $_SWF(t)$ is a measure of the cost of
establishing a fraud in a conspiracy starting at date t. The cost is
measured in the number of actions that the party perpetrating the
fraud must perform that carry a risk of exposure.

In general, the Social Work Factor will increase over time.
Perpetrating a fraud claiming that the Roman emperor Nero never
existed today would require that millions of printed histories be
erased and rewritten, every person who has ever taught or taken a
lesson in Roman history would have to participate in the fraud. The
Social Work Factor would be clearly prohibitive.

The Social Work Factor in the immediate aftermath of Nero's
assassination in 68 would have been considerably lower. While the
emperor Nero was obviously not erased from history, this did happen
to Akhenaten, an Egyptian pharaoh of the 18^th dynasty whose
monuments were dismantled, statues destroyed, and his name erased
from the lists of kings.

1.8.1. Related work

It has not escaped the notice of the author that the social work
factor might be applied as a general metric for assessing the
viability of a conspiracy hypothesis.

Applying social work factor analysis to the moon landing conspiracy
theory we note that almost all of the tens of thousands of NASA
employees who worked on the Apollo project would have had to be a
part of the conspiracy and so would an even larger number of people
who worked for NASA contractors. The cost of perpetrating the hoax
would have clearly exceeded any imaginable benefit while the risk of
the hoax being exposed would have been catastrophic.

2. The problem of trust

Traditional (symmetric key) cryptography allows two parties to
communicate securely provided they both know a particular piece of
information known as a key that must be known to encrypt or decrypt
the content. Public Key cryptography proposed by Diffie and Hellman
[Diffie76] provides much greater flexibility by using separate keys
for separate roles such that it is possible to do one without being
able to do the other. In a public key system, an encryption key
allows information to be encrypted but not to be decrypted. That
role can only be performed using the corresponding decryption key.

The Mathematical Mesh recryption services further extend the
capabilities of traditional public key infrastructures by further
partitioning of the roles associated with the private key. In the
Mesh, this capability is referred to as 'recryption' as it was
originally conceived of as being a form of Proxy Re-encryption as
described by Blaze et. al. but it might equally well be considered as
realizing distributed key generation as described by Pedersen. A
decryption key is split into two or more parts such that both parts

must be involved to complete a private key operation. These parts
are then distributed to separate parties, thus achieving
cryptographic enforcement of a separation of duties.

Public key cryptography allows many (but certainly not all)
information security concerns to be reduced to management of
cryptographic keys. If Alice knows the Bob's encryption key, she can
send Bob an encrypted message that only he can read. If Bob knows
Alice's signature key, Bob can verify that a digital signature on the
message really was created by Alice.

A Public Key Infrastructure (PKI) is a combination of technologies,
practices and services that support the management of public key
pairs. In particular, if Alice does not know Bob's public key, any
infrastructure that is designed to provide her with this information
may be regarded as a form of PKI.

The big challenge faced in the design, deployment of operation of a
PKI is that while Alice and Bob can communicate with perfect secrecy
if they use each other's actual public keys, they will have worse
than no security if an attacker can persuade them to use keys they
control instead. One of the chief concerns in PKI therefore is to
allow users to assess the level of risk they face, a quality known as
trust.

2.1. Existing approaches

Few areas of information security have engaged so much passionate
debate or diverse proposals as PKI architecture. Yet despite the
intensity of this argument the state of deployment of PKI in the
Internet has remained almost unchanged.

TLS and SSH, the only Internet security protocols that have
approached ubiquity both operate at the transport layer. The use of
IPSEC is largely limited to providing VPN access. DNSSEC remains a
work in progress. Use of end-to-end secure email messaging is
negligible and shows no sign of improvement as long as competition
between S/MIME and OpenPGP remains at a stalemate in which one has a
monopoly on mindshare and the other a monopoly on deployment.

2.1.1. Trust After First Use (TAFU)

Trust After First Use is a simple but often effective form of PKI.
Instead of trying to verify each other's public key the first time
they attempt to communicate, the parties record the public key
credentials presented in their first interaction and check that the
same credentials are presented in subsequent transactions. While
this approach does not absolutely guarantee that 'Alice' is really

talking to 'Bob', as the conversation continues over hours, months or even years, they are both assured that they are talking to the same person.

2.1.2. Direct Trust

In the direct trust model, credentials are exchanged in person. The exchange may be of the actual public key itself or by means of a 'fingerprint' which is simply a means of formatting a cryptographic digest of the key to the user.

Use of direct trust is robust and avoids the need to introduce any form of trusted third party. It is also limited for the obvious reason that it is not always possible for users to meet in person. For this reason, protocols that attempt to offer a direct trust model often turn out to be being used in trust-after-first-use mode in practice when the behavior of users is examined.

2.1.3. Certificate Authority

The archetype of what is generally considered to be 'PKI' was introduced in Kohnfelder's 1978 Msc. Thesis [Kohnfelder78] . A Certificate Authority (CA)whose signature key is known to all the participants issues certificates binding the user's public key to their name and/or contact address(es).

This approach forms the basis of almost every widely deployed PKI including the EMV PKI that support smart card payments, the CableLabs PKI that supports the use of set top boxes to access copyright protected content and the WebPKI mentioned earlier that supports the use of TLS in online commerce.

One area in which the CA model has not met with widespread success is the provision of end-to-end secure email described in the original paper. Despite the fact that S/MIME secure email has been supported by practically every major email client for over 20 years, only a small number of users are aware that email encryption is supported and even fewer user it on a regular basis.

One of the reasons for this lack of uptake is the lack of uptake itself. Until a critical mass of users is established, the network effect presents as the chicken and egg problem. Another reason for the failure is the sheer inconvenience use of S/MIME presents to the user. Obtaining, installing and maintaining certificates requires significant user effort and knowledge. But even if these obstacles are addressed (as the Mesh attempts to do), as far as the open Internet is concerned, S/MIME provides little or no benefit over a

direct trust model because there is no equivalent of the WebPKI for email.

Most CAs that operate WebPKI services also offer S/MIME PKI services, but these are seldom used except by enterprises and government agencies where certificates are usually issued for internal use only.

One of the chief difficulties in establishing a MailPKI analogous to the WebPKI is the difficulty of establishing a set of validation requirements that are cost effective to users and present a meaningful social work factor to attackers.

When VeriSign began operating the first Internet CA, two classes of email certificate were offered that have since become a de facto industry standard:

Class 1: The CA verified that the subject applying for the certificate could read email sent to the address specified in the certificate.

Class 2: The requirements of class 1 plus the requirement that the certificate be issued through a Registration Authority that had been separately determined to meet the considerably more stringent validation requirements for organizations specified in class 3 and in particular, demonstrated ownership of the corresponding domain name.

Class 2 certificates were designed to be issued by organizations to their employees and arguably present a more than adequate social work factor to prevent most forms of attack. S/MIME certificates are in daily use to secure very sensitive communications relating to very high value transactions. But this represents a niche application of what was intended to be a ubiquitous infrastructure that would eventually secure every email communication.

The only type of certificate that the typical Internet user can obtain is class 1 which at best offers a small improvement on social work factor over Trust After First Use.

2.1.4. Web of Trust

The concept of the Web of Trust was introduced by Zimmerman with the launch of PGP. It represents the antithesis of the hierarchical CA model then being proposed for the Privacy Enhance Mail scheme being considered by the IETF at the time. A core objection to this model was the fact that users could only communicate securely by obtaining a certificate from a CA. The goal of PGP was to democratize the process by making every user a trust provider.

Like S/MIME, OpenPGP protocol has achieved some measure of success
but has fallen far short of its original goal of becoming ubiquitous
and almost none of the users have participated in the Web of Trust.

One of the chief technical limitations of the Web of Trust is that
trust degrades over distance. An introduction from a friend of a
friend has less value than one from a friend. As the number of users
gets larger, the chains of trust get longer, and the trustworthiness
of the link becomes smaller.

Another limitation is that as is fitting for a concept launched at
the high tide of postmodernism, the trust provided is inherently
relative. Every user has a different view of the Web of Trust and
thus a different degree of trust in the other users. This makes it
impossible for a commercial service to offer to navigate the Web of
Trust on a user's behalf.

2.1.5. Chained notary

The rise of BitCoin [Bitcoin] and the blockchain technology on which
it is based have given rise to numerous proposals that make use of a
tamper-evident notary as either the basis for a new PKI (e.g.
NameCoin [Namecoin]) or to provide additional audit controls for an
existing PKI (e.g. Certificate Transparency [RFC6962]).

The principle of making a digital notary service tamper-evident by
means of combining each output of the notary with the input of its
successor using a cryptographic digest was proposed in 1991 by Haber
and Stornetta [Haber91] . Every output of the notary depends on every
one of the previous inputs. Thus, any attempt to modify an input
will cause every subsequent output to be invalidated.

Notaries operating according to these principles can quickly achieve
prohibitively high social work factors by simply signing their output
values at regular intervals and publishing a record of the signed
values. Any attempt by the notary to tamper with the log will
produce a non-repudiable proof of the defection. Thus once an input
value is enrolled in a chained notary, the social work factor for
modifying that input subsequent to that becomes the same as the
social work factor for subverting the notary and every party that has
a record of the signed outputs of that notary.

Enrolling the signed outputs of one notary as an input to another
independently operated notary establishes a circumstance in which it
is not possible for one notary to defect unless the other does as
well. Applying the same principle to a collection of notaries
establishes a circumstance in which it is not possible for any notary
to defect without that defection becoming evident unless every other

notary also defects. If such infrastructures are operated in
different countries by a variety of reputable notaries, the social
work factor of modifying an input after it is enrolled may be
considered as to rapidly approach infinity.

One corollary of this effect is that just as there is only one global
postal system, one telephone system and one Internet, convergence of
the chained notary infrastructure is also inevitable. Users seeking
the highest possible degree of tamper evidence will seek out notaries
that cross notarize with the widest and most diverse range of other
notaries. I propose a name for this emergent infrastructure, the
Internotary.

According to the image presented in the popular press, it is the
minting of new cryptocurrency that provides stability to the
distributed leger at the heart of BitCoin, Etherium and their many
imitators. The fact that notaries that do not require proof of work,
proof of stake or any other form of seigniorage offer the same social
work factor (effectively infinite) as those that do demonstrates that
it is not necessary to consume nation-state level quantities of
electricity to operate such infrastructures.

The attraction of employing such notaries in a PKI system is that the
social work factor to forge a credential prior to a date that has
already been notarized as past is infinite. It is obvious that
almost none of the thousands of OpenPGP keys registered with the key
server infrastructure for 'Barack Obama' are genuine and so all the
registered keys are untrustworthy. But if it was known that one
particular key had been registered in the 1980s, before Obama had
become a political leader, that particular key would be considerably
more trustworthy than the rest.

The use of chained notaries may be viewed as providing a distributed
form of Trust After First Use. The first use event in this case is
the enrollment of the event in the notary. Instead of Alice having
to engage in separate first use events with Bob, Carol, Doug and
every other user she interacts with, a single first use event with
the internotary supports all her existing and future contacts.

2.1.6. A blended approach

As we have seen, different PKI architectures have emerged to serve
different communities of use by offering different forms of trust.
The trust provided by the OpenPGP and S/MIME PKIs to the communities
they serve is distinct. The S/MIME PKI does not provide a useful
means of establishing a trusted relationship in a personal capacity.
The OpenPGP PKI is not appropriate for establishing a trust
relationship in an enterprise capacity. Yet despite this obvious

difference in capabilities, there has been no convergence between
these competing approaches in the past two decades.

The only convergence in approach that has developed over this period
is within the applications that rely on PKI. Most SSH clients and
servers make provision for use of CA issued certificates for
authentication. Most email clients may be configured to support
OpenPGP in addition to S/MIME.

While offering the choice of CA issued, direct trust or Web of Trust
credentials is better than insisting on the use of the one, true PKI,
this approach is less powerful than a blended approach allowing the
user to make use of all of them.

In the blended approach, every user is a trust provider and can
provide endorsements to other user and some (but not necessarily all)
users have CA issued certificates.

This approach follows the same patterns that have been applied in the
issue of government credentials for centuries. In many countries,
passport applications must be endorsed by either a member of a
profession that has frequent interaction with the public (e.g.
doctors, lawyers and clerics), a licensed and registered set of
public notaries or both.

Analysis of the blended approach in terms of work factor reveals the
surprising result that it can achieve a higher social work factor
than either the CA model alone or the Web of Trust model alone.

Consider the case that Alice and Bob have each obtained a certificate
that presents a Social Work Factor of $10. Applying the CA model in
isolation, $10 is the limit to the SFW that can be achieved. But if
Alice and Bob were to meet and exchange endorsements, the SFW may be
increased by up to $10. If the exchange of endorsements is made in
person by means of some QR code mediated cryptographic protocol, we
might reasonably ascribe a SWF of $20 to each credential.

This higher SWF can now be used to evaluate the value of endorsements
issued by Alice and Bob to user Carol and of Carol to Doug, neither
of whom has a CA issued certificates. While the SWF of Carol is
certainly less than $20 and the SWF or Doug is even lower, it is
certainly greater than $0.

While these particular values are given for the sake of example, it
is clearly the case that as with the WebPKI, the blended approach
permits trust to be quantified according to objective criteria even
if the reliability of the values assigned remains subjective. The
Google Page Rank algorithm did not have to be perfect to be useful

and just as the deployment of the Web spurred the development of
engines offering better and more accurate search engines, deployment
of blended PKI may be reasonably be expected to lead to the
development of better and more accurate means of evaluating trust.

The power of the blended approach is that it provides the reach of
the Web of Trust model with the resilience of the CA model while
permitting a measurable improvement in work factor over both.

Combining the blended trust model with the internotary model allows
these SWF values to be fixed in time. It is one thing for an
attacker to spend $100 to impersonate the President of the United
States. It is quite another for an attacker to spend $100 per target
on every person who might become President of the United States in 20
years' time.

3. The Mesh of Trust

The purpose of the Mathematical Mesh is to put the user rather than
the designer in control of their trust infrastructure. To this end,
the Mesh supports use of any credential issued by any form of PKI and
provides a means of using these credentials in a blended model.

3.1. Master Profile

The Mesh provides an infrastructure that enables a user to manage all
the cryptographic keys and other infrastructure that are necessary to
provide security.

A Mesh master profile is the root of trust for each user's personal
PKI. By definition, every device, every application key that is a
part of user's personal Mesh profile is ultimately authenticated
either directly or indirectly by the signature key published in the
master profile.

Unlike user keys in traditional PKIs, a Mesh master profile is
designed to permit (but not require) life long use. A Master profile
can be revoked but does not expire. It is not possible to change the
signature key in a master profile. Should a compromise occur, a new
master profile must be created.

3.2. Uniform Data Fingerprints

Direct trust in the Mesh is realized through use of Uniform Data
Fingerprints (UDF) [draft-hallambaker-mesh-udf] . A UDF consists of a
cryptographic digest (e.g. SHA-2-512) over a data sequence and a
content type identifier.

UDFs are presented as a Base32 encoded sequence with separators every
25 characters. UDFs may be presented at different precisions
according to the intended use. The 25-character presentation
provides a work factor of 2^117 and is short enough to put on a
business card or present as a QR code. The 50-character presentation
provides a work factor of 2^242 and is compact enough to be used in a
configuration file.

For example, the UDF of the text/plain sequence "UDF Data Value" may
be presented in either of the following forms:

MDDK7-N6A72-7AJZN-OSTRX-XKS7D
MDDK7-N6A72-7AJZN-OSTRX-XKS7D-JAFXI-6OZSL-U2VOA-TZQ6J-MHPTS

The UDF of a user's master profile signature key is used as a
persistent, permanent identifier of the user that is unique to them
and will remain constant for their entire life unless they have
reason to replace their master profile with a new one. The exchange
of master profile UDFs is the means by which Mesh users establish
direct trust.

3.3. Strong Internet Names

A Strong Internet name (SIN) [draft-hallambaker-mesh-udf] is a valid
Internet address that contains a UDF fingerprint of a security policy
describing interpretation of that name.

While a SIN creates a strong binding between an Internet address and
a security policy, it does not provide a mechanism for discovery of
the security policy. Nor is it necessarily the case that this is
publicly available.

For example, Example Inc holds the domain name example.com and has
deployed a private CA whose root of trust is a PKIX certificate with
the UDF fingerprint MB2GK-6DUF5-YGYYL-JNY5E-RWSHZ.

Alice is an employee of Example Inc., she uses three email addresses:

For example, Example Inc holds the domain name example.com and has
deployed a private CA whose root of trust is a PKIX certificate with
the UDF fingerprint MB2GK-6DUF5-YGYYL-JNY5E-RWSHZ.

Alice is an employee of Example Inc., she uses three email addresses:

alice@example.com A regular email address (not a SIN).

alice@mm--mb2gk-6duf5-ygyyl-jny5e-rwshz.example.com A strong email
   address that is backwards compatible.

alice@example.com.mm--mb2gk-6duf5-ygyyl-jny5e-rwshz A strong email
   address that is backwards incompatible.

Use of SINs allows the use of a direct trust model to provide end-to-
end security using existing, unmodified email clients and other
Internet applications.

For example, Bob might use Microsoft Outlook 2019, an email
application that has no support for SINs as his email client. He
configures Outlook to direct outbound mail through a SIN-aware proxy
service. When Bob attempts to send mail to a strong email address
for Alice, the proxy recognizes that the email address is a SIN and
ensures that the necessary security enhancements are applied to meet
the implicit security policy.

3.4. Trust notary

A Mesh trust notary is a chained notary service that accepts
notarization requests from users and enrolls them in a publicly
visible, tamper-evident, append-only log.

The practices for operation of the trust notary are currently
undefined but should be expected to follow the approach described
above.

The trust notary protocol provides support for establishing an
internotary through cross certification. The append only log format
is a DARE Container [draft-hallambaker-mesh-dare] , the service
protocol is currently in development.

3.5. Endorsement

An endorsement is a document submitted to a trust notary that
includes a claim of the form 'public key X is held by user Y'. Mesh
endorsements may be issued by CAs or by ordinary users.

3.6. Evaluating trust

One of the chief advantages of the World Wide Web over previous
networked hypertext proposals was that it provided no means of
searching for content. While the lack of a search capability was an
obstacle to content discovery in the early Web, competing solutions
to meeting this need were deployed, revised and replaced.

The Mesh takes the same approach to evaluation of trust. The Mesh
provides an infrastructure for expression of trust claims but is
silent on their interpretation. As with the development of search

for the Web, the evaluation of trust in the Mesh is left to the
application of venture capital to deep AI.

[4](). Conclusions

This paper describes the principal approaches used to establish
Internet trust, a means of evaluating them and a proposed successor.
It now remains to determine the effectiveness of the proposed
approach by attempting deployment.

[5](). Security Considerations

This document describes the means by which interparty identification
risk is managed and controlled in the Mathematical Mesh.

The security considerations for use and implementation of Mesh
services and applications are described in the Mesh Security
Considerations guide [draft-hallambaker-mesh-security] .

[6](). Acknowledgements

A list of people who have contributed to the design of the Mesh is
presented in [draft-hallambaker-mesh-architecture] .

[7](). References

[7.1](). Normative References

[draft-hallambaker-mesh-architecture]
          Hallam-Baker, P., "Mathematical Mesh 3.0 Part I:
          Architecture Guide", draft-hallambaker-mesh-
          architecture-08 (work in progress), July 2019.

[draft-hallambaker-mesh-security]
          Hallam-Baker, P., "Mathematical Mesh Part VII: Security
          Considerations", draft-hallambaker-mesh-security-00 (work
          in progress), April 2019.

[7.2](). Informative References

[Adrian2015]
          Adrian, D., "Weak Diffie-Hellman and the Logjam Attack",
          October 2015.

[Bitcoin] Finley, K., "After 10 Years, Bitcoin Has Changed
          Everything?And Nothing", November 2018.

[Diffie76]
          Diffie, W. and M. Hellman, "New Directions in
          Cryptography", November 1976.

[draft-hallambaker-mesh-dare]
          Hallam-Baker, P., "Mathematical Mesh 3.0 Part III : Data
          At Rest Encryption (DARE)", draft-hallambaker-mesh-dare-02
          (work in progress), July 2019.

[draft-hallambaker-mesh-udf]
          Hallam-Baker, P., "Mathematical Mesh 3.0 Part II: Uniform
          Data Fingerprint.", draft-hallambaker-mesh-udf-03 (work in
          progress), July 2019.

[Haber91] Haber, S. and W. Stornetta, "How to Time-Stamp a Digital
          Document", 1991.

[Intel2018]
          Bell, L., "Intel delays 10nm Cannon Lake processors,
          again, until late 2019", July 2018.

[Kohnfelder78]
          Kohnfelder, L., "Towards a Practical Public-Key
          Cryptosystem", May 1978.

[Namecoin]
          Inc., N., "Namecoin Web Site", 2019.

[Ratchet] Marlinspike, M. and T. Perrin, "The Double Ratchet
          Algorithm", November 2016.

[RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand
          Key Derivation Function (HKDF)", RFC 5869,
          DOI 10.17487/RFC5869, May 2010.

[RFC6246] Sajassi, A., Brockners, F., Mohan, D., and Y. Serbest,
          "Virtual Private LAN Service (VPLS) Interoperability with
          Customer Edge (CE) Bridges", RFC 6246,
          DOI 10.17487/RFC6246, June 2011.

[RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate
          Transparency", RFC 6962, DOI 10.17487/RFC6962, June 2013.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol
          Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018.

[Schneier2013]
          Schneier, B., "Defending Against Crypto Backdoors",
          October 2013.

[Shannon1949]
          Shannon, C., "Communication Theory of Secrecy Systems",
          1949.

7.3. URIs

   [1] http://mathmesh.com/Documents/draft-hallambaker-mesh-trust.html

Author's Address

   Phillip Hallam-Baker

   Email: phill@hallambaker.com