

Outline of a proposal responding to E.U. and U.S. calls for trustworthy global-scale IdM and CKM designs

Benjamin Gittins
Synaptic Laboratories Limited
PO Box 5, Nadur, Gozo, NDR-1000, Malta, Europe
cto@pqs.io

ABSTRACT

In 2007, the E.U. FP6 SecurIST called [31] for trustworthy international identity management (**IdM**) that was user-centric. In 2009, the U.S. Department of Homeland Security (**DHS**) called [28] for trustworthy [70] global-scale IdM and the U.S. National Institute of Standards and Technology (**NIST**) called [13] for new cryptographic key management (**CKM**) designs. In this paper we outline the core architecture for (apparently) the first globally scalable, post quantum secure, symmetric key based *platform* for provisioning IdM, key distribution/agreement and inter-enterprise CKM services. Our proposal employs a decentralised trust model that exploits compartmentalisation, redundancy and diversification simultaneously across service provider, software developer, hardware vendor, class of cryptographic primitive, and protocol axis. It employs behavioural analysis techniques and supports the collaborative management of international name spaces, management of client transactions using public identifiers and supports user-centric cross-cutting control mechanisms. Our proposal is suitable for use with commercial off the shelf hardware and is designed to wrap-around and protect the output of existing security deployments. The platform addresses the U.S. Networking and Information Technology Research and Development Program (**NITRD**) call [56] to create a digital immune system (multi-layered protection, decentralised control, diversity, pattern recognition), the DHS call [28] for combating insider attacks and malware, achieving survivability and availability, and NIST managers' call for a CKM design supporting billions of users without the use of public key technologies [13]. This proposal has been designed as part of our Trustworthy Resilient Universal Secure Infrastructure Platform project [38].

Categories and Subject Descriptors

E.3 [Data encryption]; C.2.1 [Computer-communications networks]: Network architecture and design—*distributed networks, store and forward networks, network topology*.

Permission to make digital or hard copies of all or part of this work is granted provided the copies bear this notice and the full citation on the first page. Version 1.1 as published on ePrint (March 14, 2011).

This work is based on an earlier work: Overview of SLL's proposal in response to NIST's call for new global IdM-CKM designs without Public Keys, in Proceedings of the 6th Annual Workshop on Cyber Security and Information Intelligence Research (April 21-23, 2010) © ACM, 2010.

1. INTRODUCTION

In 1976, two fundamentally different techniques were published that enabled authenticated private conversations between any two parties over a public network. The first unnamed technique, proposed by W. Diffie, M. Hellman and L. Lamport, employed a symmetric key distribution protocol [30] exploiting m key distribution nodes (aka key distribution centers) [10] that was secure against a collusion of up to $m-1$ participating key distribution nodes. We name this proposal **DHL-SKD**. The second technique, proposed by W. Diffie, M. Hellman and R. Merkle, employed public key encryption and required digital signatures [29]. Unfortunately, derivatives [21] of Shor's 1994 quantum algorithm [65] threaten the confidentiality and integrity of classical public key algorithms [55] based on the problem of factoring large numbers, the discrete logarithm problem, or elliptic curve schemes. Many identity based encryption schemes [18] are based on the same problems and so are also at risk. Identification of a trustworthy post quantum secure asymmetric key exchange remains an open hard problem [16], [60]. Independent of the quantum computing threat there are many other serious flaws [40], [46], [47] that have plagued the civilian global-scale PKI and fundamentally undermine its utility [48].

In 2009 the U.S. President's cyberspace policy review [70] near term action plan called for game-changing technologies that have the potential to enhance the security, reliability and trustworthiness of digital infrastructure and to "*build a cybersecurity-based identity management vision that addresses privacy and civil liberties interests*". The DHS responded to this call with their "Roadmap for Cybersecurity Research" [28] which outlines 11 current hard problems in information security, including global-scale IdM. NIST formally responded to the policy review by declaring that the development of new CKM capable of billions of users must be part of the U.S. national cybersecurity initiatives [13]. In both cases, current technologies are not considered adequate.

In this paper we show how to extend the 1976 symmetric key distribution scheme [30] to create a platform for a semi-online global-scale IdM, key distribution/agreement and inter-enterprise CKM that responds to the above calls. The fundamental principles of our design were well received by J. Patarin and L. Goubin in their 2008 review. The precursor to this paper was peer-reviewed and published by the 2010 CSIIRW-6 [35]. The applicability of our model in

network behavioural analysis (and remote malware detection) was published [50] by O. McCusker and others at the NATO IA&CDS [6]. Network behavioural extensions to our model were also published at ORNL CSIIRW-6 [49]. Our design was published at the 2010 IEEE Key Management Summit [36], [37].

2. STRUCTURE OF THE PAPER

This paper has 2 parts: the context around our proposal and the proposal itself.

Part 1: Context. In §3 we re-evaluate the original watershed decision that promoted public key distribution over symmetric techniques [30]. In §4 we survey the drivers motivating our work: In §4.1 we outline design requirements found in the ‘Spirit of Laws’ political theory treatise [27]. In §4.2 we summarise E.U. FP6 SecurIST’s published position on user centricity. In §4.3 we recite the 11 current hard challenges to achieving trustworthiness as identified by the U.S. DHS. Finally in §4.4 we outline NIST’s 2009 CKM drivers [13]. In §5 we observe that IdMS and CKMS are interdependent §5.1 and discuss trustworthiness framed in the context of global-scale IdM-CKM §5.2. The cryptographic foundations of our platform rely on symmetric techniques §6: In §6.1 we perform a short survey of early symmetric key distribution results. In §6.2 we quote W. Diffie, M. Hellman, and L. Lamport’s description of their symmetric key distribution proposal and make observations on it in §6.3.

Part 2: Proposal. In §7 we describe the network topology of our IdM-CKM proposal: In §7.1 we rewrite the 1976 DHL-SKD [§6.2] to scale wrt. service providers and server nodes. In §7.2 we describe the network topology between a client and a store and forward node (SFN) and in §7.3 the network topology between a pair of SFN. In §7.4 we compare the security properties of homogenous and diversified realisations of this topology. In §7.5 we illustrate the network connectivity between two clients on the network. Finally in §7.6 we indicate various deployment strategies and walk through a pedagogical global-scale deployment scenario involving communication between regional and international clients. In §8 we make explicit our a priori vulnerability assumptions §8.1 and we survey the design’s conformance with our drivers in §8.2. In §9 we sketch how to provision a variety of services. In §9.1 we describe how high-availability communications is achieved between nodes of the cryptographic overlay network (IdM-CKM platform). In §9.2 we sketch how to assigning public identifiers to clients. In §9.3 we outline the context of inter-enterprise key management, describe how to scale secret sharing schemes wrt. to the number of shares, and then apply this within the context of our global-scale cryptographic overlay network (IdM-CKM platform). In §9.4 we describe how clients recall keys on demand. In §9.5 we describe the push based distribution of keys. By using (or rewriting) §9.4 and §9.5 we show how to perform: key agreement in §9.6, key agreement with crypto diversity in §9.7, provision authenticated assertion records in §9.8, provision secure file sharing in §9.9 and provision secure messaging in §9.10. In §9.11 we sketch how ExoskeletonsTM (protocol aware point-to-point tunnels) employing services provisioned by our proposed IdM-CKM platform can protect deployed infrastructure without requiring changes to software or hardware im-

plementations of standards based security standards. In §10 we discuss (dis)trust and accountability before ending with a conclusion in §11.

Part 1: Context

3. RE-EVALUATING PKI DRIVERS

In 1976, W. Diffie and M. Hellman (D&H) conjectured [29, 30] that offline public key infrastructure (PKI) was required to achieve scalability and availability. Today online techniques are routinely applied to scale offline X.509 based PKI. This negation prompts us to reconsider their drivers.

Driver 1: Avoid secure key distribution channels.

The use of self-signed certificates relaxed the original requirement for a trusted courier to deliver pair-wise unique symmetric keys down to the authenticated delivery of a public root certificate. The mass availability of CPU based smart cards is relatively new phenomena that was unavailable to D&H in 1976. These programmable smart card modules, when mounted on reels, can be efficiently used as a secure distribution channel for pair-wise unique symmetric keys. An enrolling party can visually fingerprint the smart card modules (using high-resolution laser imaging) and install custom applets before supplying them to service providers for key-injection operations. The tokens can then be optically inspected (for similarity and tampering) and electronically queried on return. The enrolling parties then act as authenticated distribution channels by supplying smart cards to end-users. Public key techniques, using merkle tree digital signature based algorithms [51, 25], can *also* be used to validate the authenticity of the smart cards.

Driver 2: Enable private conversations between any two parties regardless of whether they have ever communicated before.

In 1976, D&H held that offline public key distribution was more bandwidth/latency efficient at key distribution than their $m-1$ secure symmetric key distribution proposal (DHL-SKD). Today, public key distribution with Online Certificate Status Protocol (OCSP) involves a network transaction. In 1976, ARPANET [26] and X.25 [22] clients were not designed to support concurrent network sessions. Today, concurrency is uniformly available which reduces the network transaction latency by a factor of m . Today, the difference in network latencies between public key distribution with OCSP checking and DHL-SKD is much less than anticipated in 1976. With the advent of CPU based smart cards, DHL-SKD network costs can be amortised by securely managing symmetric keys over multiple network sessions and by performing key derivation.

Driver 3: Enable scalable authentication of communication parties.

In 1976, D&H expressed concern with node scalability and network availability issues and sought offline methods. Offline authentication operations in X.509 [41] require certificates and digital signature technologies. The responsibility for certificate/public key life-cycle management (discovery, validation) was shifted away from online servers. Users were left to find their own ad hoc solutions. Today, this heavy burden shifted to users is considered a serious hindrance to ubiquitous encryption [56]. These problems do not exist in symmetric systems. In key distribu-

tion and key translation architectures [10] pair-wise unique symmetric keys are employed to perform mutual authentication and key exchanges with low CPU overhead, either directly or through tickets. Advantageously, all reachable identities are discoverable in one location and the freshest key material is always supplied to users.

Driver 4: Remove the need for online servers. Summarising 3 results from P. Gutmann’s paper [40]: 1) It is not possible to explicitly *validate* certificates in the X.509, instead offline certificate *revocation* lists are used. 2) The Online Certificate Status Protocol (**OCSP**) is a proxy service designed to improve the scalability of the certificate *revocation* lists. 3) The OCSP requires computationally expensive digital signatures for authenticated operations. OCSP also has vulnerabilities [46].

4. DRIVERS MOTIVATING OUR WORK

We propose that cryptographic systems should seek to address relevant requirements and calls as found below.

4.1 L’esprit des lois design requirements

The “Spirit of Laws” is a treatise on political theory first published anonymously by Charles de Secondat, Baron de Montesquieu in 1748 [27]. Montesquieu was the most frequently quoted authority on government and politics in colonial pre-revolutionary British America, cited more by the American founders than any source except for the Bible [45]. Montesquieu advocated constitutionalism, the separation of powers, checks and balances, the preservation of civil liberties, and the rule of law with the objective to reduce citizens fear of the political system. The important role true anonymity (as opposed to Government revocable pseudo-anonymity) has played historically in democracies should be considered in the design of, and laws concerning, IdM and CKM systems.

4.2 E.U. FP6 SecurIST on user centrality

Based on text and quotes from SecurIST publications [62, 31]: “*In the E.U., privacy is generally defined as a right of self-determination, namely, the right of individuals to determine for themselves when, how and to what extent information about them is communicated to others.*” SecurIST calls for international user-centric IdM in which the end users are empowered to determine his or her own security and dependability requirements and preferences. “*User-centric mechanisms are required to allow controlled release of personal, preference-related and location-based information, and to deliver assurances to owners about how personal information will be used by third parties.*” This marks a shift “*from Security and Dependability by 20th century central command and control approaches*”, towards architectures that could lead to an “*open and trustworthy Information Society through empowerment*” of the individual with the purpose of protecting the central systems, the citizen and society interests (i.e. protecting the legitimate interests of all stake holders). “*Responsibility, authority and control have to move more towards the end user.*”

4.3 U.S. DHS on trustworthiness

The Nov. 2009 DHS “Roadmap for Cybersecurity Research” [28] outlines 11 current hard problems, eight of which “*were*

selected as the hardest and most critical challenges that must be addressed by information security research if trustworthy systems envisioned by the U.S. Government are to be built.” The 8 challenges being: global-scale IdM, insider threats, availability of time-critical systems, building scalable secure systems, situational understanding and attack attribution, information provenance, privacy aware security and enterprise-level security metrics. The remaining 3 hard challenges being: system-evaluation life cycle, usable security and combating malware and botnets. Information processing systems striving for trustworthiness should address as many of these challenges they can from the onset of their design. The call for global-scale IdM was stressed again in June 2010 [33].

4.4 U.S. NIST’s CKM drivers

At the 2009 NIST Cryptographic Key Management (CKM) Workshop [13], NIST managers identified that new CKM designs should be highly available, fault tolerant, secure against destructive attacks, scalable to billions of users, enable the ubiquitous take up of encryption, be secure against quantum computer attacks and use means other than public key technologies. Additionally they must support accountability, auditing, policy management, and be interoperable. NIST subsequently published their draft “Framework for Designing Cryptographic Key Management Systems” (SP 800-130) [14] in June 2010 resulting in comments received [17]. Over 90% of the points raised in NIST’s summary of public feedback comments [24] presented at the second NIST CKM Workshop [3] were submitted by Synaptic Laboratories. Among other things, our feedback identified the need for the CKMS framework to be reconciled with other standards, special publications, guidance and forms such as SP 800-57 [12], the DHS Cybersecurity Roadmap [28], IEC 61508 Safety Integrity Levels [4], the US National Strategy for Trusted Identities in Cyberspace [7] and so on. At the end of the 2-day workshop, the results of the 2 breakout study groups correlated with our recommendations.

5. GLOBAL-SCALE IDM AND CKM

5.1 IdMS and CKMS are interdependent

The New Oxford American Dictionary defines a secret as “*something that is kept or meant to be kept unknown or unseen by others*”. Cryptographic systems employ a) CKMS to manage keys and establish authenticated private channels and b) IdMS to identify and authenticate identities. Electronic IdMS use cryptography to authenticate identities and physical IdMS to identify people. We can’t define an electronic-IdMS without defining a CKMS and vice versa. IdMS and CKMS are as interdependent as Yin and Yang. Global-scale cryptographic systems require collaboration between CKM, electronic IdM and physical IdM specialists.

5.2 Trustworthy global-scale IdM-CKM

To paraphrase Montesquieu, a global-scale IdM-CKM should be set up so no stake-holder need be afraid of another. This requires a conceptual shift away from the ‘us vs. them’ adversarial model inherited from the military origins of cryptography and towards an inclusive regulative system between peers. We assert that principles and requirements outlined in §1 and §4 can be embodied and realised in a unified trustworthy and cost-effective IdM-CKM system. A system




Figure 1: The Yin-Yang of IdMS and CKMS

that enhances democratic principles and protects the legitimate and diversified interests of all stake holders/users, even in a global context of competing nation-states. A global-scale IdM-CKM system provides the opportunity to realise user-centricity envisioned by the E.U. and others in a way not possible with today's uncoordinated silo'd (federated) based security solutions. In this paper we outline the core architecture of a global-scale platform that can be extended to comprehensively address international CKM, electronic-IDM and physical-IDM in a co-ordinated but distributed, decentralised and diversified manner. Our proposal exploits diversity in membership to improve security through a system of checks-and-balances and separation of powers in a way that ensures the system remains highly available and robust to all stake holders. Diversity used in this manner also encourages international competition in the open market place.

The IdM-CKM proposal as described in this paper protects clients from security compromises as a result of latent vulnerabilities or malware present in the software or hardware used by IdM-CKM service providers, or by the service provider's privileged technical or managerial staff. Our IdM-CKM proposal will achieve further improved confidentiality, integrity and availability properties for the IdM-CKM service providers when the IdM-CKM server software is hosted on our Trustworthy Resilient Universal Secure Infrastructure Platform proposal [38].

6. SYMMETRIC KEY DISTRIBUTION (SKD)

6.1 A short survey of early SKD results

In 1970 H. Feistel [32] described the use of symmetric keys to perform mutual authentication and this was applied to a network context by D. Branstad in 1973 [19] and 1975 [20]. In 1976 W. Diffie, M. Hellman and L. Lamport proposed the use of m key distribution nodes, where $m \geq 2$ [30]. We call this unnamed proposal **DHL-SKD**. S. Kent's 1976 thesis [42] gave the first description of a cryptographic system that employed two factor authentication, $m \geq 1$ symmetric key distribution networks, chaining of symmetric secrets between network sessions (stored on magnetic cards), and the authenticated encryption of data. Our proposal extends these results.

6.2 The DHL-SKD proposal

With reference to figure 6.2 we quote [30]: “A small number m of the network's nodes will function as ‘key distribution

nodes’. Each user has m keys, one for communicating with each of these m nodes. These keys vary from user to user, so while each user must remember only m keys, each of the key distribution nodes remembers n , one for each user of the net. When users A and B wish to establish a secure connection they contact the m key distribution nodes and receive one randomly chosen key from each. These keys are sent in encrypted form using the keys which the users share with the respective nodes. Upon receiving these keys, the conversants each compute the exclusive or of the m keys received to obtain a single key which is then used to secure a private conversation. None of the nodes involved can violate this privacy individually. Only if all m nodes are compromised will the security of this connection fail.” The paper goes on to say under the usual idealized security assumptions DHL-SKD is secure against a collusion of any combination of $m-1$ key distribution nodes. If one or more of the key distribution nodes is performing a denial of service attack the users select a subset n of the m key distribution nodes, in which case the protocol is secure against of any combination of $(n-1)$ key distribution nodes.

6.3 Our observations on DHL-SKD

The 1976 DHL-SKD proposal did not specify if the m key distribution nodes are operated by 1 or m different service providers (that is, did we achieve m -Independence). It did not specify if the m key distribution nodes should run on identical platforms or exploit hardware and/or software diversity [23, 58]. The DHL-SKD proposal can be implemented using NIST FIPS 140 approved symmetric cryptographic primitives/modes of operation. NIST Advanced Encryption Standard [53] with 256-bit keys is widely considered post quantum secure for encryption [39]. NIST Secure Hash Algorithm [54] with 256-bit digest is widely considered PQS for message authentication. Key distribution nodes (**KDN**) and key translation centers (**KTC**) are both a type of secure store-and-forward node (**SFN**). For the purpose of two devices establishing a secure authenticated network connection, each of the m key distribution nodes in DHL-SKD can be trivially adapted to operate as key translation centre without invalidating the original security argument. An *idealised* key translation centre (a link-level secure key relay service) can be rewritten as: a network of unsecured processing elements enclosed and operating within the protection of a TEMPEST SDIP-27 [1] certified electromagnetic shielded enclosure performing key translation operations.

Part 2: Proposal

7. SLL'S IDM-CKM TOPOLOGY




Figure 2: Topology of DHL-SKD with $m = 3$ key distribution nodes and 2 clients A and B




Figure 3: Topology of our scalable architecture

Our primary contribution in this paper is to outline a global-scale cryptographic overlay network, derivable from the DHL-SKD proposal. This overlay network is a platform suitable for delivering a wide range of inter-organisation, authenticated, policy driven, store-and-forward based cryptographic services such as secure messaging, key distribution, key agreement, key storage, and IdM operations. Our cryptographic overlay network has a semi-regular topology with certain well defined topological constraints that ensure consistent operational performance. Similar to the DHL-SKD model, most client transactions provisioned from a IdM-CKM deployment are distributed across a subset $x \geq 3$ service providers that the client is enrolled with, those service providers selected from x of the c confederations. In some cases a client may perform administrative operations, such as billing, with a single store and forward node.

7.1 Rewriting DHL-SKD to scale

With reference to figure 3 we consider each deployment of the DHL-SKD scheme to be an instance of a cryptographic overlay network (i.e. there may be multiple independent deployments of the DHL-SKD scheme). We substitute the m idealised key distribution nodes (KDN) of the DHL-SKD proposal with c confederations (illustrated as pentagons), where $c = m$. Each of the c confederations has at least 1 service provider (illustrated as a diamond). A service provider is assigned exclusively to one of the c confederations in this cryptographic overlay network instance. (A service provider may participate simultaneously in multiple cryptographic overlay network deployments.) Each service provider must have at least 1 store and forward node (SFN) (octagon). In practice, each confederation should have at least x SFN, where $x \geq 3$. Each of the x SFN shares at least one pairwise unique symmetric key (≥ 256 -bits in length) with the other $x-1$ SFN in its confederation. Each of the x SFN operates within the protection of an TEMPEST SDIP-27 electromagnetic shielded enclosure [1, 2]. Each of the x SFN communicate with the other $x-1$ SFN in a confederation using post quantum secure authenticated encrypted communications using the corresponding symmetric key. Efficient methods of $m-1$ post quantum secure bootstrapping of confederations and incrementally enrolling SFN are known.

In this way we have rewritten a SFN implemented as a network of unsecured processing elements enclosed within a single TEMPEST SDIP-27 certified electromagnetic shielded enclosure as a SFN implemented as a network of TEMPEST




Figure 4: Topology between a client and a SFN

certified processing elements, where the TEMPEST certified processing elements communicate with each other using pairwise unique post quantum secure channels. Under idealised conditions, both versions of the SFN description are at least post quantum secure against outside adversaries.

7.2 Topology between a client and a SFN

Each client is enrolled with c store and forward nodes (SFN), one from each of the c confederations. Figure 4 illustrates the topology between one client and one of the c SFN. In this higher assurance embodiment the client has two CPU based smart cards Si and So and the SFN has two network attached hardware security modules Ho and Hi . The smart card Si and Hi share a pairwise unique symmetric key (≥ 256 -bits in length). Likewise, the smart card So and Ho share a pairwise unique symmetric key (≥ 256 -bits in length). The hardware security module Hi executes the SFN server logic, and the smart card Si executes the SFN client logic. The hardware security module Ho and So execute point-to-point secure tunnel logic. The encrypted ciphertext generated in the inner communications layer between Si and Hi is re-encrypted by So and Ho resulting in an outer layer of security. Each enrolled client has a total of $2c$ pairwise unique symmetric keys. Key injection operations, performed by c service providers, can be executed while the smart cards processors are still on reels. Each SFN has 2 pairwise unique symmetric keys with every enrolled client. Ideally, Ho and Hi operate within the protection of a TEMPEST certified enclosure, and So and Si employ side-channel and fault injection protection mechanisms.

7.3 Topology between a pair of SFN

In preferred higher assurance embodiments, each of the x SFN shares at least two pairwise unique symmetric key (≥ 256 -bits in length) with the other $x-1$ SFN in its confederation. Communications between every pair of SFN involves an inner and outer layer of communications security, similar to the technique described in §7.2. Cross-cutting communication between SFN may also be required. These pair-wise unique keys would be exchanged online, on demand, as required.

7.4 Assigning agents to the abstract topology

The security properties of our proposal vary based on the agents participating.

Homogeneity: Let us consider a small **degenerate** homogenous deployment with $c = 4$ confederations, 1 service provider per confederation, and 1 store and forward node (SFN) per service provider. We assign all these resources to one division of one organisation. The $2c$ hardware security modules are provided by the same hardware security module vendor. The $2c$ modules are installed and run from the same

room. The operations of the inner and outer smart cards are all assigned to one smart card. All smart cards enrolled into the system are from the same smart card vendor. The protocol software for the hardware security modules and smart cards is implemented by one software developer. The deployment standardises entirely on NIST standards running in identical modes of operation (AES-CTR, SHA2-HMAC) for all cryptographic operations. In this way the hypothetical degenerate deployment strives to aggregate control and responsibility towards fewer agents, making the system more vulnerable to common mode of failures.

Diversity: Let us consider a similar sized deployment which preferentially exploits diversity and independence. It has $c = 4$ confederations, 1 service provider per confederation, and 1 store and forward node (SFN) per service provider. For simplicity of description, we select only two different smart card vendors, a first vendor for Si and a second vendor for So . For simplicity of description, all clients enrolled into the system will use a token from the same 2 vendors. We assign each of 4 confederations one of the following countries {Iceland, Russia, China, United States}. The 4 service providers are autonomous/independent organisations (wrt. other service providers) and each service provider is incorporated in the country assigned to their respective confederation. The 4 SFN are installed in the country of their respective confederation. Each of the 4 SFN are randomly assigned 2 different hardware security module vendors from the set of all available hardware security module vendors, where that random selection is refined to ensure each hardware security vendor is present within the deployment and also well represented (avoid heavy biases). Each service provider assigns one of their divisions to vetting/implementing their local copy of the software for the smart card So and hardware security module Ho for their instance of the outer layer. (In this way, the client smart card So receives c applets implementing the outer layer operations, a different applet for each SFN.) Each service provider assigns a different division of their organisation to vetting their local copy of the software used for their instance of the inner layer software for the smart card Si and hardware security module Hi . The inner layer employs NIST standards based cryptographic primitives and modes of operation. The outer layer employs alternate cryptographic primitives, such as non-US regional standards such as the GOST standards [5] and [71] for the Russian provider or other popular primitives. In this way the preferred hypothetical deployment strives for diversity, separation of powers (influence) in a redundant way with the aims of improving security (and at times improving availability).

7.5 Enrolled clients

Figure 5 illustrates 3 confederations of a IdM-CKM overlay network deployment. Label A illustrates a first client that is enrolled with three store and forward nodes (SFN) selected from the three confederations. Label B illustrates a second client that is enrolled with three store and forward nodes (SFN) selected from the same three confederations. Recall that every SFN shares a pair-wise unique symmetric key with every other SFN in a confederation, permitting a post quantum secure channel between every pair of SFN in that confederation. Client A and Client B can establish post quantum secure link-level encrypted paths across each




Figure 5: Paths between 2 clients over 3 confed.

confederation of the overlay network.

7.6 Deployment strategies and example

A deployment of the IdM-CKM overlay network can organise its confederations in a variety of ways including: as a global system with service providers grouped by aligned countries, as a regional system, as a national system with service providers grouped by different agencies/organisations, or even as an enterprise system. Deployments of the IdM-CKM overlay network can be layered, permitting a global IdM-CKM infrastructure for international communications, and several independent regional, national, industry controlled overlay network for localised traffic as described below.

Hypothetical global scale deployment. One possible configuration of a global scale IdM-CKM overlay network has $c = 6$ confederations with membership criteria as follows: { {UK, USA, CA, AU, NZ}, {EU member states (excluding the UK which is already assigned)}, {Arab States}, {Asian States}, {African States}, {all other remaining States} }. (Other configurations may be more desirable). Each confederation has 4 service providers, and each service provider has 4 SFN. The deployment employs diversification techniques as described in §7.4.

One of many possible international deployment layers participating in the global scale deployment. The five countries in the first confederation of the global scale deployment can reuse their existing investments and simultaneously participate in a second IdM-CKM overlay network with $c = 5$ confederations with membership criteria as follows: {UK, USA, CA, AU, NZ}. This configuration may be highly desirable for supporting their inter-government communications, and for commercially sensitive transactions between those countries.

Case use of an enrolled client. A client from Canada may be enrolled in their choice of service providers, one from each of the following countries: {UK, USA, CA, Iceland, Dubai, South Korea}. That Canadian client can use these 6 service providers to enrol (exchange keys) online with other service providers participating as clients in the global system. In this way the Canadian client can enrol with both a AU and a NZ service provider. The Canadian client exchanging key material with a New Zealand client may choose to set a default policy to use service providers from the ‘Five-Eyes’ countries {UK, USA, CA, AU, NZ} of which they are both a member, thereby minimising information leak-

age to other countries that may not normally share intelligence with this group of Nation-states. The same Canadian client, exchanging key materials with a Norwegian client would most likely use the service providers selected from the global-scale deployment as this guarantees availability of secure paths for exchanging key material. The ability for the client to chose their preferred service providers, and which service providers to use depending on the transaction, supports the U.S. NITRD's call for Tailored Trustworthy Spaces, and the E.U. call for user-centric empowerment.

Metadata. The regrettable leakage of connection information in our proposal is comparable with the information leakage already resulting from the international use of secure socket layer/transport layer security. SSL/TLS network communications over the Internet between clients in different countries leaks information to the countries that the Internet packet traverses (and any countries the respective certificate revocation query traverses), and through those countries to their respective allies they share intelligence with, and so on. e.g. SSL connections between two Japanese citizens located in their own country potentially leaks information to American intelligence organisations if they rely on the U.S. based Verisign as one of their root certificate authorities. Our proposal, as described above, can reduce this type of leakage.

8. ASSUMPTIONS AND CONFORMANCE

8.1 A priori vulnerability assumptions

In our proposal we work under the conservative assumption that latent unknown security vulnerabilities (malicious or otherwise) are *present* within the software and/or hardware of a cryptographic overlay network deployment. Our design objective is to limit one or more colluding agents to induce a service failure wrt. availability, confidentiality, integrity or maintainability of the system. Our goal is to ensure the reliability and safety of operations on behalf of all stake holders/users, even in the face of destructive attacks/natural disasters.

8.2 Conformance with our drivers

We will now comment on the conformance of our design with our drivers.

Achieve scalability of topology. Our IdM-CKM proposal permits scalability in the number of confederations, the number of different service providers within a confederation, and the number of store and forward nodes (SFN) within a service provider. Uniform performance characteristics across the system can be met by defining quality of service level requirements that must be met by service providers with regard to every client token they manage. This permits variation in the aggregate computing power of service providers and even confederations (on the provision that the number of confederations in a deployment is larger than the number of confederations each token is enrolled in).

Achieve scalability of provisioned services. This requires care in selecting what services to offer, and how to deliver them. Advantageously the constraints behind the semi-regular topology of the IdM-CKM permit certain assumptions and design optimisations to be made. For example, the number of SFN within a confederation required to

forward messages in a client-to-client transaction is upper bound to 2. This property ensures certain security properties are present, and that wide area network latencies are upper bound even as the system scales. In our experience so far, essential cross-cutting services can be efficiently realised in an arbitrarily scalable manner. The mappings of Uniform Resource Identifiers [15] to the SFNs managing the clients' tokens associated with that identifier change infrequently and can be synchronised system-wide relatively easily. In contrast, a volatile database mapping of client tokens with their current Internet protocol address is comparatively burdensome and entirely unnecessary because online tokens can disclose their SFN directly between each other over unsecured network channels (zero overhead for the SFN), and validated at the end of the cryptographic handshake.

Redundancy. IdM-CKM platform employs client transaction redundancy across confederations, and client communication security redundancy through layering of independently keyed cryptographic protocols (preferably with different cryptographic primitives).

Diversity. Our IdM-CKM platform employs diversity [23, 58] at every point of redundancy in the design, including diversification across confederations, service providers, software and hardware vendors, class of cryptographic primitives, and in layers of communication security.

Separation of powers. Separation of powers is where the functions of a system are divided into separate and independent powers and areas of responsibility. Similar to the application of separation of powers within a country, this principle is applied within the context of a service provider in our design - typically only one organisation is assigned to each component. In the same way that we can observe redundancy with diversity when we look at two or more countries that both implement separation of powers, we also see redundancy and diversity at a system-wide deployment level of our platform. Where this property has limited benefit on a day-to-day basis for citizens in the context of the organisation of nation-states, in our case every client gains improved assurances on each transaction they perform. By applying diversity at every point of redundancy in our model, we limit the total amount of power/authority/control/influence a vendor or component has within a cryptographic overlay network deployment.

Checks and balances. As (almost) all client transactions are distributed redundantly across several autonomous service providers there is implicitly some form of checks and balances in place for those transactions. This property is made explicit through cross-cutting negotiation between participating service providers, and possibly one or more other representative authorities, to determine if the requested client transaction is authorised.

Multilayered protection: Our IdM-CKM design promotes layering of different secure communication protocols for both client-to-store and forward node, and client to client operations (see §9.11). In addition we propose services provisioned by the platform implement behavioural analysis techniques that employ human-in-the-loop techniques to mitigate misconduct by users and privileged administrators.

Decentralised control: The core of our IdM-CKM platform is decentralised organisation of (semi-)autonomous service providers that collaborate together to perform client transactions. In an international deployment, there is no system-wide single point of authority/control. Furthermore, the layering of communication security protocols ensures that the protocols employed within a deployed system are not under any one organisations control.

Useability: Our IdM-CKM platform employs smart cards to simplify client side key-management. The ability to globally co-ordinate the assignment of public Uniform Resource Identifiers with clients, in an online system that ensures freshness of key material and validation of identifiers permits vastly simplified key management over current X.509 type solutions.

Collaborative management of name spaces: A single global-scale deployment of our IdM-CKM platform can act as a clearing house for each nation's registers (assertion providers) for people (registry of births, deaths, and marriages), corporations (corporate registry) and top-level domain names (.com, .br, .fr, ...). Each client can consult with the service providers it has a relationship with from the *c* different confederations to form a consensus opinion on the validity of an assertion, without the client having to know (or have a relationship with) the internationally recognised authority for the different types of assertions. Additional assertion providers may be responsible for managing assertions made from a specific portion of a name space (IANA¹, au.IANA, com.au.IANA, compay.com.au.IANA), for assigning tokens to identity assertions, for creating and assigning roles and responsibilities within an organisation, and so on.

User centricity: Each person and organisation is a single logical entity, independent of the ability for a person to have multiple names and roles or an organisation to have multiple directors and authorised agents within it. In a global-scale IdM-CKM deployment with multiple assertion providers attesting various attributes regarding the existence and status (e.g. dead or alive) of an entity, and the mapping of a token to that entity (or authorised agent for that entity), it becomes possible to provide a cross-cutting user-centric view of the information managed by a global-scale IdM-CKM system. This can be done while simultaneously ensuring that every organisation managing a relationship with that entity has a 'per organisation unique identifier' (pseudonym).

Privacy enhancing technology: Services provisioned from our IdM-CKM platform can be privacy enhancing in the way that is envisioned by the EU STORK [67] and US NSTIC [69] initiatives. e.g. ensuring conditional release of information and the use of pseudonyms where desired/required. As we proceed to advance the design we will be looking for opportunities to minimise the amount of meta-data trivially leaked to service providers. We will be asking questions such as: can a service provider manage data in a user-centric cross-cutting way while masking these relationships from the service providers through the use of indirection/pseudonyms and further compartmentalisation of information.

Achieve fault-tolerance: Redundancy can be employed within the compute and storage elements of each store and forward node to improve availability of services in the case of hardware faults. The use of distributed atomic transaction (begin, commit, rollback) based programming techniques by a service provider can be used to mask hardware failures of store and forward nodes without exposing the hardware failure to clients [11]. The presence of redundant service providers can be adapted to increase availability of the system in the case of the failure of a service provider. This may vary depending on the client transaction. With client-to-client key exchange, fault-tolerance is achieved through negotiation between the clients by allowing the number of participating service providers to be reduced in response to an unavailable/misbehaving service provider. With the remote storage of data across multiple service providers the encoding of client data using an all-or-nothing transformation [59] that is further encoded with parity and then distributed over multiple service providers permits the client to access their data even if one service provider is unavailable.

Achieve availability: The presence of fault-tolerance in a design leads to improved service availability. In our design, where a service provider has one or more store and forward nodes, it is possible to dynamically re-assign the store and forward nodes responsible for processing token requests in response to work-load within that service-provider. More uniform assignment of work load increases the responsiveness and availability of the service. The systematic application of quality-of-service techniques through the system can increase the availability of mission-critical services and permit price differentiation of services.

Combat insider attacks: Some high-availability systems achieve software [44] and/or hardware fault tolerance [11] using redundant implementations of the same function, running on independent circuits, potentially implemented by different teams where the output of the functions passes through a ballot monitor. The principle of redundancy and diversification has been adapted to create intrusion tolerant systems such as SITAR [72] where it is assumed an adversary can introduce a service failure in software executing within the system. Our IdM-CKM platform when implemented with diversity can mitigate a wide variety of insider attacks residing within the hardware or software of any component from compromising a client's security. In contrast to some intrusion tolerant systems which seek to detect and respond to intrusion events on an otherwise un-compromised deployment, our design explicitly assumes the intruder has a persistent presence inside the deployment and seeks to limit their ability to leverage that presence against a IdM-CKM client. This line of approach to combating insider attacks has been refined further in our Trustworthy Resilient Universal Secure Infrastructure Platform [38].

Survivability against destructive attacks: Physically destructive attacks resulting from natural disasters or deliberate malicious human acts can result in catastrophic service failure at a site. If an attack is experienced by a service provider at one site, continuity of services for that service provider is possible if redundant systems are available at one or more physically different sites. If one service provider experiences total catastrophic service failure, it is

¹Internet Assigned Numbers Authority (IANA)

possible for clients to negotiate relationships with other service providers and restore redundancy in any information stored in the IdM-CKM deployments by substitution operations performed by the failed service provider with a new service provider.

Situational awareness: Unlike X.509 PKI systems which are intentionally designed as predominantly offline systems, (semi-)online IdM-CKM systems are designed to actively participate in the delivery of many client transactions. Online systems can be trivially adapted to maintain state, and this state can be used to achieve situational awareness. For example, online IdM-CKM systems can *selectively* store information about the access patterns of a Client, or an IP address. In this way our IdM-CKM platform can support situational awareness and provide useful and appropriate services to clients.

For a *service provider-client* relationship to be trustworthy (e.g. doctor-patient, attorney-client, specialist-layman, computer-user, cloud-user, ...) the party entrusted with sensitive information must not exploit that information in a way that undermines the legitimate interests of that stakeholder. Likewise, trustworthy information processing systems (human or automated) should be designed to minimise the amount of exploitable clear-text information they receive, while ensuring they leverage sensitive clear text information entrusted to them solely for the benefit of the client (virtue). Systems that (individually or systematically) violate this axiomatic principle undermine the community and cannot/will not be trusted by the same. e.g. A corporation of lawyers would irrevocably undermine their client's trust if they exposed sensitive personally information. Likewise, it follows that to realise a global-scale trustworthy IdM-CKM deployment, as is called for by the E.U. and U.S. Government, it must be virtuous and uphold this axiom.

Behavioural analysis and pattern recognition: Behavioural analysis techniques can be used to detect behaviours which may indicate possible security risks. To maintain user centricity, behavioural analysis should be performed for the benefit of each stakeholder in the system. Each stakeholder may have their own unique behavioural analysis policies which the system should enforce. A range of default policies should also be made available to make these services immediately available. A human-in-the-loop process should be used to manage risk events detected by the system. A client should be able to delegate the human-in-the-loop to their outsourced managed security solution provider if they desire.

Combating malware and botnets: U.S. Sonalysts Inc is designing a distributed sensor system for the Internet (Oculex), which delivers policy-driven behavioral-based trust of hosts, derived from analysing aggregated network behaviors over multiple time scales for threat behaviors. Malware and botnets often exhibit distinctive behaviors that can be remotely detected by sensor networks. Behavioral analysis of sensor data, when done without identity, enables the sharing of actionable information without infringing upon the privacy of individuals or the community. On the remote detection of certain classes of malware, notification (via reverse look-up through the IdM-CKM platform) can then lead to

remedial action to the relevant stakeholders. Separation of powers should be enforced, ensuring that identity information is not supplied from the IdM-CKM to the sensor network, and behavioural data exchanged between sensor nodes should not be supplied to the IdM-CKM deployment. See our co-authored paper for more information [49].

Post quantum secure: Our IdM-CKM platform relies entirely on symmetric cryptographic primitives which can select operational parameters (such as key length and digest length) that are widely considered to be both classically and post quantum secure. These primitives are available and widely trusted today.

9. SERVICE PROVISIONING

Our proposed IdM-CKM platform can be used to provision a wide range of cryptographic services. In this section we outline how communication is achieved with high availability within the IdM-CKM overlay network and then outline several cryptographic client services.

9.1 Overlay network communications

Most client transactions and all client-to-client transactions provisioned by a IdM-CKM deployment are distributed across a subset $x \geq 3$ service providers that the client is enrolled with, those service providers being selected from x of the c confederations. In some cases a client may perform administrative operations, such as billing, with a single service provider.

With reference to figure 6, in preferred high availability embodiments clients are enrolled with two store and forward nodes (**SFN**) owned by the same service provider paired in an {active, hot standby} buddy system. The hot standby node is illustrated as a light grey octagon with thick black border. The pairing is on a per-client basis. The SFN buddies may be physically located in two geographically separated sites (located in the east and west borders of a country or continent). The client has a pairwise unique symmetric key with each SFN (the enrolment with the hot standby SFN may be performed online with first use). If the active SFN becomes unavailable, the client continues the transaction on the hot standby SFN (which becomes the active SFN). The client may be directed by a service provider to exchange one of the SFN pairs with a different SFN managed by the same service provider in response to work-load balancing or hardware failure. The low-level details of how the buddy system should be implemented is outside the scope of this paper. In our model a client can establish an authenticated secure channel with every (active or hot standby) SFN it is enrolled with as described in §7.2. A client can request a first active SFN it is enrolled with to establish a secure connection with a second active SFN in the same confederation. The preferred secure connection between the first and second SFN is described in §7.3. In high-availability deployments the hot-standby SFN mirror the active SFN, as illustrated in figure 6. A reference to a SFN now implies the active SFN unless otherwise indicated.

The client can relay messages through the first SFN to the second SFN, and through the second SFN to any of the clients enrolled with the second SFN. In this case, the first SFN is responsible for identifying the client to the second




Figure 6: 2 clients, SFN buddy system, 3 confed.

SFN, and the second SFN is responsible for forwarding that identification information to other clients. Clients are responsible for correlating message parts from across the x confederations and checking that the identity assertions are the same from the x independent service providers.

In this paper, each SFN can reliably store data (with associated use policies), on behalf of an enrolled client in non-volatile memory. In high-availability systems, this data is replicated across the buddy system.

9.2 Assigning public identifiers to clients

After a client's token(s) are enrolled with a service provider it can be assigned one or more public identifiers (such as universal resource identifiers [15]) at low cost using an automated challenge-response process establishing the token's user has control of an e-mail account/website. This process can be reinforced through manual checking of physical credentials when higher levels of assurance are required. The process for the transfer and control of an identifier varies depending on the level of attestation previously provided and it's description is outside the scope of this paper. The redundancy in validating identifiers helps protect name spaces as assets of their respective owners/stake holders.

9.3 Inter-enterprise key management

Context. In the context of managing the private key of root certificate authorities, some commercial enterprise CKM products offer M of N split key controls where ($2 \leq M \leq 5$), ($2 \leq N \leq 7$) and ($M \leq N$). The value of the secret is split into N shares, where any M combination of those N shares can reconstruct the original value of the secret. Split-key schemes are also known as secret sharing schemes [63]. The N person controls are often managed by people employed by one enterprise/organisation. All transactions performed with that split key requires the participation of M agents. Transactions include key exchanges, message signing, changing the membership of N and so on. The logistic effort to perform a transaction increases as the value of M increases. In practice, split-knowledge schemes do not arbitrarily scale wrt. the number of shares.

In the context of managing enterprise keys, sometimes M of N split authentication access controls are used. In this case the full value of the secret is entrusted to a hardware security module. The hardware security module is supplied a policy that requires M out of N parties to authorise a transaction on that secret. The stake-holders have to trust the hardware security module to consistently enforce that pol-

icy. Adding and removing authorised parties is easier in this case, as knowledge of the secret is not split across the N parties. To improve system availability and transaction workload capacity in conventional enterprise CKM deployments, two or more hardware security modules may mirror each others configuration. Unfortunately simple replication in this way increases the risk of a single hardware security module failure compromising that deployment. Furthermore, the attractiveness of attacking a hardware security module tends to increase along with the number of stake-holders that are dependent on it.

In an inter-enterprise key exchange environment, if the hardware security module is under sole control of one organisation, dependent organisations may have little to no assurances regarding their ability to control and audit transactions. If the full value of the key is known to that HSM, then a dependent organisation may have no assurances with regard their ability to control who can discover the value of the key. If one or more of those hardware security modules is attached to the Internet, the stake-holders require additional assurances that the split-authentication access controls cannot be subverted remotely. Unfortunately, it does not appear possible for any single vendor to demonstrate to their clients that this type of vulnerability is not present in their device². One or more malicious software developers may covertly install vulnerabilities that could be exploited. Likewise, vendors of hardware security modules may be compelled to (covertly) install kill-switches or interception technologies in hardware security modules intended for local and/or foreign markets. Back-doors may be lurking in the components that hardware security vendors employ in their hardware security modules. Countries such as America are extremely concerned regarding the possibility of back-doors and kill switches under foreign control [8]. To quote the E.U. SecurIST [31]: “*The lack of trust is one of the main barriers for the establishment of a secure and dependable Information Society.*”

To summarise, it is not possible in practice to arbitrarily scale the number of parties that share partial knowledge of a secret, and schemes where knowledge of the value of the secret is not split across multiple parties are limited in the level of security assurances they can offer to stake-holders.

Scaling split key operations. We are proposing an inter-enterprise IdM-CKM scheme where the knowledge of a secret is split over a small manageable number of shares $3 \leq c \leq 7$ and the authentication and access control is managed independently for each share (resulting in scalable split-authentication controls). Given it is unreasonable to require a vendor/organisation to demonstrate the complete absence of vulnerabilities in a product/process, we propose that each of the c shares is managed by a different service provider while ensuring those c shares are managed by several different hardware security module vendors, ideally c different vendors. See §7.4 and §7.6 for more information on our preferred deployment strategy.

²To quote B. Schneier: “*No one can guarantee 100% security*” ... “*There’s no test possible that can prove the **absence** of flaws.*” ... “*A good cryptographic system strikes a balance between what is possible and what is acceptable.*” [61]




Figure 7: Inter-enterprise key storage

In this configuration clients rely on other service providers to independently manage some or all of the shares. If a secret is encoded in a N of M split key scheme where $M = c$ and $N = M - 1$ the client's secret remains *available* if one of the M service provider becomes unavailable and the clients secret remains *secure* against a collusion (or simultaneous compromise) of $N - 1$ service providers.

If a client wants to increase their security assurances they can participate as a service provider for their own transactions. This capability is available to each client, and between clients. i.e. two clients can be actively participating in the management of key shares between themselves (this requires their smart cards to be enrolled with each others hardware security modules).

Pedagogical example. Long-lived key material with associated access policies can be stored and enforced by a global-scale IdM-CKM deployment. In our model we require that client transactions provisioned from a IdM-CKM deployment are distributed across a subset $x \geq 3$ service providers (that the client is enrolled with) from x different confederations.

With reference to figure 7 client A receives the value of a 256-bit secret it wishes to store and an access use policy associated with that secret. In this pedagogical illustration client A is enrolled with $c = 3$ service providers from c confederations. (In preferred embodiments $4 \leq c \leq 7$). Client A wants to encode and distribute the value of the secret across the c confederations ensuring no service provider can discover the value of the secret, while ensuring the secret can be reconstructed if any 1 of the c service providers is unavailable (a $c - 1$ of c secret sharing scheme). In our illustration, one of the service providers will be assigned the parity of the encoded secret value distributed over the other $c - 1$ service providers. To ensure $\approx 2^{255}$ security against brute-force guessing attacks by one of the participating service providers, we require each share to be at least 256-bits in length. Client A allocates a contiguous array of 64 bytes (512 bits) in length in protected memory. Client A stores {the value of the 256-bit secret, the value of a 256-bit nonce} in the array. Client A encodes the 64-byte array using an unkeyed all-or-nothing transformation [59]. The 64-byte en-

coded message is partitioned into 2 shares of 256-bits in length ($share_1, share_2$). A 256-bit parity is created by calculating $share_3 = share_1 \oplus share_2$, where $\oplus$ is a binary exclusive OR operation operating on a 256-bit word. It is possible to reconstruct the original secret from any combination of 2 of the 3 shares in the usual way. Client A creates a meta-data record associate with each of the shares. The meta-data for each share stores information about the key, including the key type, key length, how many shares the key has been split into, which share of the split key is managed by this meta-data, and who is permitted to access the key. See [14] for a detailed description on the recommended fields required for key meta-data. Client A securely sends the first share and it's associated meta-data to it's enrolled service provider in the first confederation and receives back a public identifier. Client A repeats this process for the second and third shares to the second and third service providers in the second and third confederations respectively. Client A assembles the three public identifiers into a composite public identifier for that key material. The service providers are entrusted to enforce that policy with regard to their share of the key.

9.4 On demand recall of keys

Continuing from the previous paragraph and with reference to figure 7, client A has now encoded and stored a 256-bit secret across the c confederations. The secret is stored with associated meta-data instructing the store and forward nodes (SFN) how to manage that key material.

In a traditional 'enterprise CKM solution' key material is requested on-demand by one or more clients listed in the associated access policy. To achieve this operation in our design client A makes the composite public identifier for the key material known to client B. To access the key material, client B establishes secure authenticated connections with the c SFN it is enrolled with. Client B sends a request to it's c SFN to access the key material associated with the composite public identifier. Each of client B's SFN independently evaluate the request, establish a secure connection with the SFN assigned to managing the key material in their confederation, and forward the request. At this point the SFN responsible for storing the key material for Client A have all received an authenticated request from a SFN within their confederation. A cross-cutting query is performed by Client A's c SFN to establish that they have all received the same request (checks and balances). Having established a consensus to perform the request, each of client A's c SFN securely forward their respective share of the key material, along with associated policies to client B's c SFN. This material is then securely relayed to client B. Client B now has sufficient information to authenticate that the key material is from client A and to reconstruct the original value of the key. This process can be readily adapted to support a range of key management operations as authorised by the meta-data associated with that key material.

9.5 Push based distribution of keys

Continuing from the previous paragraph, the meta-data associated with client A's key material could instruct client A's store and forward nodes (SFN) to notify one or more target clients that key material is available. Client A's SFN are then responsible for identifying the SFN associate with

each of the target clients and informing them of the key's availability. Target clients are notified of the availability of key material immediately, or when they next log-in with their SFN. The policy for that key material may optionally instruct client A's SFN to delete the key material after a) all targets have successfully received the key material, b) an expiration time, c) or both.

9.6 Key agreement

Both client A and client B can use the push based distribution of keys to securely exchange nonce. Client A and client B would each receive the other's nonce via the IdM-CKM deployment, concatenate the 2 nonce in the same order and supply the output of the concatenation operation as input to a cryptographic hash function, using the resulting digest as shared key material.

9.7 Key agreement with crypto diversity

In a two-pass online key agreement protocol that exploits symmetric and asymmetric technologies, client A and client B use the push based key distribution function to securely exchange their respective public keys in the first pass. Client A and client B receive the authenticated public key of the other client. In the second pass client A and client B use the push based key distribution function to securely exchange the ciphertext resulting from their respective public key encryption of a nonce. Client A and B receive the ciphertext, decrypt the nonces, concatenate the 2 nonce in the same order and supply the output of the concatenation operation as input to a cryptographic hash function, using the resulting digest as shared key material. Advantageously, this method protects the ciphertext of public key operations using post quantum secure symmetric techniques and depending on the strength of the asymmetric algorithm chosen it may provide additional protection from a collusion of all participating service providers. Client A and Client B can choose to use different asymmetric algorithms to further increase crypto diversity or to satisfy their respective regional security standards.

9.8 Assertion records

Instead of key material, the client-to-client key distribution techniques described in §9.4 can be used to store public (or private) authenticated assertions such as SAML assertions [57], domain name server resource records [52], resource permissions, and so on.

9.9 Secure file systems

Instead of key material, the client-to-client key distribution techniques described in §9.4 can be used to store long-lived objects, the objects being either directories or files.

9.10 Secure messaging

Instead of key material, the client-to-client key distribution techniques described in §9.5 can be used to transmit short-lived objects, such as instant messages or e-mails.

9.11 Protecting deployed infrastructure

Secure tunnels are designed to wrap around and protect the (potentially insecure) output of programs without changing them. Protocol aware secure tunnels, which we call

ExoskeletonsTM, would provide improved post quantum secure protection for the output of each network session generated by implementations of at-risk public key dependent security standards such as SSL/TLS, IPsec, RADIUS, and SSH. This capability could protect today's massive classically secure PKI deployments in a non-disruptive manner. Exoskeletons can be developed in a controlled environment without requiring existing standards to be adjusted. The technology can then be incrementally or rapidly deployed on a moments notice as desired/required.

10. (DIS)TRUST AND ACCOUNTABILITY

It is not appropriate to design global systems where insiders must be trusted. Today, approximately 86% of fraud happens by management level staff against their own organisation, in part **because they can** circumvent security mechanisms intended to prevent fraud [43]. Global systems that centralize trust in one 'trusted third party' (TTP) fuel the risk of cyber fraud and cyber war because they require users to absolutely trust the integrity of that trusted third party (or in the case of PKI, some 20+ Root certificate authorities [66], [40]).

Security systems should be designed so that no stakeholder is in fear of another. This can be done by redundantly distributing the execution of each provisioned service across m autonomously owned/managed service providers to mitigate insider fraud/attacks. Users do not need to buy into the altruism of any service provider. Instead users may choose to place their confidence in the mutual distrust and/or competitiveness between service providers. Such systems already employ "separation of powers" and can be adapted to employ cross-cutting "checks and balances" [27], provide redundant transaction audit logs to all users of the system, prevent liability shifting [9], and provide balanced security, accountability and privacy [68] for all stakeholders/users [31], [64].

11. CONCLUSION

Federal agencies and co-ordinating bodies in the U.S. and E.U. are calling in unison for trustworthy, resilient and dependable information and communications infrastructure that protects civil liberties and is user-centric. Calls for new trustworthy international/global-scale identity management and cryptographic key management designs have been made. This paper is a response to those calls. We have introduced (apparently) the first globally scalable, symmetric, IdM-CKM platform that is robust against a wide range of insider attacks. We have listed the ways our proposal addresses several drivers identified by U.S. and E.U. authorities. Our architecture can be derived from an existing proposal [30] which is already considered post quantum secure. Our proposal is practical, cost effective and can be implemented using commercial off-the-shelf hardware and implemented using NIST (or regional standards based) symmetric ciphers and hash functions which are already accepted to be post quantum secure. Our proposal can be used to provision a diverse range of client services by mapping traditionally specialised services (key distribution, key agreement, key management, name server, assertion server, file server, secure email, secure instant messaging) in a uniform way onto a authenticated store-and-forward network that exploits compartmentalisation, redundancy and diversification throughout the design. Our proposal can be used to protect exist-

ing at-risk public key cryptosystems. Our feedback [34] to the NIST draft framework for designing CKMS [14] appears to have also been positively received. Our (internationally distributed) decentralised trust model employs the democracy supporting Principles Of Laws and can be deployed in a manner that empowers all stake-holders and promotes goodwill and engenders trust between nations.

12. REFERENCES

- [1] Compromising emanations laboratory test standard. SECAN Doctrine and Information Publication SDIP-27 Level A, NATO.
- [2] Laboratory test standard for protected facility equipment. SECAN Doctrine and Information Publication SDIP-27 Level B, NATO.
- [3] Cryptographic Key Management Workshop 2010. Project, National Institute of Standards and Technology, Sep. 2010.
- [4] Functional safety of electrical/electronic/programmable electronic safety-related systems. IEC 61508, International Electrotechnical Commission, 2010.
- [5] GOST 28147-89: Encryption, Decryption, and Message Authentication Code (MAC) Algorithms. In V. Dolmatov, editor, *RFC5830*, United States, Mar. 2010. RFC Editor.
- [6] NATO Information Assurance and Cyber Defence Symposium in Turkey. Website, Apr. 2010.
- [7] National Strategy for Trusted Identities in Cyberspace. Project, National Institute of Standards and Technology, January 2011.
- [8] S. Adele. The Hunt for the Kill Switch. In *IEEE Spectrum*. IEEE, May 2008.
- [9] R. J. Anderson. Liability and computer security: Nine principles. In *ESORICS '94: Proceedings of the Third European Symposium on Research in Computer Security*, volume 875 of *LNCS*, pages 231–245, London, UK, Nov. 1994. Springer-Verlag.
- [10] ANSI. Financial institution key management (wholesale). Technical report, American National Standards Institute, 1985.
- [11] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. In *IEEE Transactions on dependable and secure computing*, volume 1, Jan. 2004.
- [12] E. Barker, W. Barker, W. Burr, W. Polk, and M. Smid. Recommendations for key management – part 1: General (revised). Special Publication 800-57 Part 1, National Institute of Standards and Technology, Mar. 2007.
- [13] E. Barker, D. Branstad, S. Chokhani, and M. Smid. Cryptographic key management workshop summary (final). Interagency Report 7609, National Institute of Standards and Technology, June 2009.
- [14] E. Barker, D. Branstad, S. Chokhani, and M. Smid. A Framework for Designing Cryptographic Key Management Systems. (Draft) Special Publication 800-130, National Institute of Standards and Technology, June 2010.
- [15] T. Berners-Lee, R. Feilding, and L. Masinter. Uniform Resource Identifier (URI): Generic Syntax. In *RFC3986*, United States, Jan. 2005. RFC Editor.
- [16] D. J. Bernstein, T. Lange, and P.-L. Cayrel. Post-quantum cryptography. Website, July 2009. Available at <http://www.pqcrypto.org>.
- [17] V. Bharadwaj, I. Clover, S. Eddy, B. Gittins, B. Nixon, S. Saha, and C.-R. Tsai. Comments Received on SP 800-130. Comments, National Institute of Standards and Technology, Aug. 2010.
- [18] D. Boneh and M. K. Franklin. Identity-based encryption from the weil pairing. In *CRYPTO '01: Proceedings of the 21st Annual International Cryptology Conference on Advances in Cryptology*, volume 2139 of *LNCS*, pages 213–229, London, UK, Aug. 2001. Springer-Verlag.
- [19] D. Branstad. Security aspects of computer networks. In *AIAA Computer Network Systems Conference, Paper 73-427*, Huntsville, Alabama, Apr. 1973. AIAA.
- [20] D. K. Branstad. Encryption protection in computer data communications. In *Proceedings Fourth Data Communications Symposium*, pages 7–9, Quebec City, Oct 1975. IEEE Computer Society.
- [21] M. Brown. Classical cryptosystems in a quantum setting. Master of mathematics in combinatorics and optimisation, Waterloo, Ontario, Canada, Apr. 2004.
- [22] CCITT. Recommendation X.25. Standard, International Telegraph and Telephone Consultative Committee, 1976.
- [23] L. Chen and A. Avizienis. N-version programming : A fault-tolerance approach to reliability of software operation. In *FTCS-8*, pages 3–9. IEEE, 1978.
- [24] S. Chokhani and M. Smid. A Summary of Public Comments on Draft Cryptographic Key Management Framework. In *NIST Key Management Workshop*. National Institute of Standards and Technology, Sep. 2010.
- [25] C. Coronado. *Provably secure and practical signature schemes*. Doctoral thesis (elib.tu-darmstadt.de/diss/000642), Technische Universität Darmstadt, Nov. 2005.
- [26] S. Crocker. Host software. In *RFC1*, United States, Sep. 1969. RFC Editor.
- [27] B. d. M. de Secondat, Charles. *The Spirit of the Laws (Originally published anonymously in 1748)*. Crowder, Wark, and Payne, 1777.
- [28] Department of Homeland Security. A Roadmap for Cybersecurity Research. Roadmap, DHS Science and Technology Directorate, Nov. 2009.
- [29] W. Diffie and M. Hellman. New directions in cryptography. In *IEEE Transactions on Information Theory*, volume 22, issue: 6, pages 644– 654, Nov. 1976.
- [30] W. Diffie and M. E. Hellman. Multiuser cryptographic techniques. In *AFIPS '76: Proceedings of the June 7-10, 1976, national computer conference and exposition*, pages 109–112, New York, NY, USA, June 1976. ACM.
- [31] Z. Dooly, J. Clarke, W. Fitzgerald, W. Donnelly, M. Riguide, and K. Howker. ICT Security and Dependability Research beyond 2010 - Final strategy. Deliverable 3.3, SecurIST EU-FP6-004547, Jan. 2007.
- [32] H. Feistel. Cryptographic coding for data bank

- privacy. Research Report RC2827, IBM T.J. Watson Res. Ctr, Yorktown Heights, N.Y., Mar. 1970.
- [33] GAO. CYBERSECURITY: Key Challenges Need to Be Addressed to Improve Research and Development. Report to Congressional Requesters GAO-10-466, United States Government Accountability Office, June 2010.
 - [34] B. Gittins. Feedback to NIST DRAFT Special Publication 800-130. Comment, Synaptic Laboratories Limited, Jun 2010.
 - [35] B. Gittins. Overview of SLL's proposal in response to NIST's call for new global IdM/CKM designs without public keys. In *Proceedings of the Sixth Annual Workshop on Cyber Security and Information Intelligence Research*, CSIIRW '10, pages 60:1–60:4, New York, NY, USA, 2010. ACM.
 - [36] B. Gittins and R. Kelson. A survey and low-level comparison of network based symmetric key distribution architectures. Video. In *IEEE Key Management Summit 2010 website*, Lake Tahoe, Nevada on May 4-5, 2010., May 2010. IEEE.
 - [37] B. Gittins and R. Kelson. Overview of SLL's proposal in response to NIST's call for new global IdM/CKM designs without PKC. Video. In *IEEE Key Management Summit 2010 website*, Lake Tahoe, Nevada on May 4-5, 2010., May 2010. IEEE.
 - [38] B. Gittins and R. Kelson. Trustworthy Resilient Universal Secure Infrastructure Platform (TruSIP) Project. Website page., ICT Gozo Malta, Jan. 2011. <https://www.ictgozomalta.eu/vision-and-projects/project-trusip-ict-ics.html>.
 - [39] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th annual ACM symposium on Theory of Computing*, Annual ACM Symposium on Theory of Computing, pages 212–219. ACM, 1996.
 - [40] P. Gutmann. *Engineering Security*. (draft book), Dec. 2009. Available at <http://www.cs.auckland.ac.nz/~pgut001/pubs/book.pdf>.
 - [41] ISO/IEC. "The Directory" series of standards. In *ISO/IEC 9594-x*.
 - [42] S. T. Kent. Encryption-based protection protocols for interactive user-computer communication. Master's thesis, Massachusetts Institute of Technology Cambridge Lab for Computer Science, May 1976.
 - [43] KPMG. Profile of a fraudster survey 2007. Forensic advisory, KPMG International, Apr. 2007.
 - [44] P. Kumar and A. P. Singh. Analyzing Software Fault-Tolerance in Real-time Systems using voting technique. Research report, University school of information technology, Sep. 2010.
 - [45] D. S. Lutz. The Relative Influence of European Writers on Late Eighteenth-Century American Political Thought. volume 78, No. 1 of *The American Political Science Review*, pages 189–197, March 1984.
 - [46] M. Marlinspike. Defeating OSCP With The Character '3'. Technical report, July 2009. Available at <http://www.thoughtcrime.org/papers/ocsp-attack.pdf>.
 - [47] M. Marlinspike. Null Prefix Attacks Against SSL/TLS Certificates. Technical report, July 2009. Available at <http://www.thoughtcrime.org/papers/null-prefix-attacks.pdf>.
 - [48] L. Martin, P. Gutmann, and B. Gittins. Is PKI really that bad? A series of postings, June 2010.
 - [49] O. McCusker, B. Gittins, J. Glanfield, S. Brunza, and S. Brooks. The Need to Consider Both Object Identity and Behavior in Establishing the Trustworthiness of Network Devices within a Smart Grid. In *Proceedings of the Sixth Annual Workshop on Cyber Security and Information Intelligence Research*, CSIIRW '10, pages 53:1–53:4, New York, NY, USA, 2010. ACM.
 - [50] O. McCusker, J. Glanfield, S. Brunza, D. C. Gates, D. J. Hugh, and D. Paterson. Combining Trust and Behavioral Analysis to Detect Security Threats in Open Environments. In *NATO IACDS 2010, RTO-MP-IST-091*, April 2010.
 - [51] R. C. Merkle. *Secrecy, authentication, and public key systems*. PhD thesis, Stanford University, CA, USA, June 1979.
 - [52] P. Mockapetris. Domain names - implementation and specification. In *RFC1510*, United States, Nov. 1987. RFC Editor.
 - [53] NIST. The advanced encryption standard. Federal Information Processing Standard 197, National Institute of Standards and Technology, Nov. 2001.
 - [54] NIST. Secure hash standard. Federal Information Processing Standard 180-2, National Institute of Standards and Technology, Aug. 2002.
 - [55] R. A. Perlner and D. A. Cooper. Quantum resistant public key cryptography: a survey. In *IDtrust '09: Proceedings of the 8th Symposium on Identity and Trust on the Internet*, volume 373 of *IDtrust*, pages 85–93, New York, NY, USA, Apr. 2009. ACM.
 - [56] QinetiQ. National Cyber Leap Year Summit 2009 – Co-Chairs' Report. On behalf of the US NITRD Program, Sep. 2009.
 - [57] N. Ragouzis, J. Hughes, R. Philpott, E. Maler, P. Madsen, and T. Scavo. Security Assertion Markup Language (SAML) V2.0 Technical Overview. OASIS Committee Draft sstc-saml-tech-overview-2.0-cd-02, Mar. 2008.
 - [58] B. Randell and J. Xu. *Software Fault Tolerance*, volume 3 of *Trends in Software*. John Wiley & Sons Ltd, 1995.
 - [59] R. Rivest. All-or-nothing encryption and the package transform. In *Fast Software Encryption*, volume 1267 of *LNCS*, pages 210–218, Jan. 1997.
 - [60] A. Rostovtsev and A. Stolbunov. Public-key cryptosystem based on isogenies. Report 2006/145, Cryptology ePrint Archive, May 2006.
 - [61] B. Schneier. Why Cryptography Is Harder Than It Looks. Essay 037, Counterpane Systems, Mar. 1997.
 - [62] SecurIST Advisory Board. Recommendations for a Security and Dependability Research Framework: from Security and Dependability by Central Command and Control to Security and Dependability by Empowerment. Deliverable 3.0, SecurIST EU-FP6-004547, Jan. 2007.
 - [63] A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, Nov. 1979.
 - [64] F. T. Sheldon, R. K. Abercrombie, and A. Mili. Methodology for evaluating security controls based on

- key performance indicators and stakeholder mission. In *HICSS '09: Proceedings of the 42nd Hawaii International Conference on System Sciences*, pages 1–10, Washington, DC, USA, Jan. 2009. IEEE Computer Society.
- [65] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, Oct. 1997.
 - [66] M. Stevens, A. Sotirov, J. Appelbaum, A. Lenstra, D. Molnar, D. A. Osvik, and B. M. M. de Weger. Short Chosen-Prefix Collisions for MD5 and the Creation of a Rogue CA Certificate. In *CRYPTO '09: Proceedings of the 29th Annual International Cryptology Conference on Advances in Cryptology*, volume 5677 of *LNCS*, pages 55–69, Berlin, Heidelberg, Aug. 2009. Springer-Verlag.
 - [67] STORK. Secure identity across borders linked. E.U. co-funded project INFISO-ICT-PSP-224993. Available at www.eid-stork.eu.
 - [68] K. Sullivan. On the anonymity “versus” accountability debate. Whitepaper, Think-Trust EU-FP7-216890, June 2010.
 - [69] US DHS and others. DRAFT National Strategy for Trusted Identities in Cyberspace. Technical report, United States Department of Homeland Security, June 2010.
 - [70] USOWH. Cyberspace policy review: Assuring a trusted and resilient information and communications infrastructure (may 26, 2009). United States, Office of the White House, May 2009.
 - [71] E. V. Dolmatov. GOST R 34.11-94: Hash Function Algorithm. In *RFC5831*, United States, Mar. 2010. RFC Editor.
 - [72] F. Wang, F. Gong, R. Sargor, K. Goseva-popstojanova, K. Trivedi, and F. Jou. SITAR: A Scalable Intrusion-Tolerant Architecture for Distributed Services. In *DARPA Information Survivability Conference and Exposition*, volume 2, pages 153 – 155, April 2003.