Identity and Access Management in a De-perimeterized World

The Missing Piece – Credential Management

Richard Moulds and Nicko van Someren
nCipher Corporation Ltd.,
Jupiter House, Station Road, Cambridge, CB2 2JD,UK
rmoulds@ncipher.com

June 2005

Abstract: Life in a de-perimeterized world means surviving out in the open in the full glare of everyone else. Safety relies on the ability to verify who you trust and share information with them confidentially. A precursor to this is the ability to identify – everything. The current Identity and Access Management (IAM) market is in a state of chaos, lacking focus and pursuing the goals of vendors rather than users. This should be of great concern to those subscribing to the Jericho vision[1] as failure to deliver a secure and unified means to validate identity and issue other credentials and rights will act as a barrier to achieving the many benefits of de-perimeterization.

This paper makes an assessment of the IAM market, identifies issues that are not currently addressed, describes in outline a credential management solution and recommends priorities for further work within the Jericho forum.

1 Introduction

The IAM market is receiving considerable interest because it claims to address a number of the issues on the 'to do list' of many IT managers and also because it is good at making its own news. After three years of almost continual consolidation involving billions of dollars, you would be excused for thinking that there is a race on. The problem is that it's not clear what constitutes the finish line or even over what ground the race is being run.

IAM refers to a class of functionality and systems that in principle integrate and interoperate. The analyst firm Gartner has produced a useful mapping of these components which is shown below in figure 1.

[1] Visioning White Paper, Jericho Forum – see http://www.opengroup.org/projects/jericho/uploads/40/6809/vision_wp.pdf

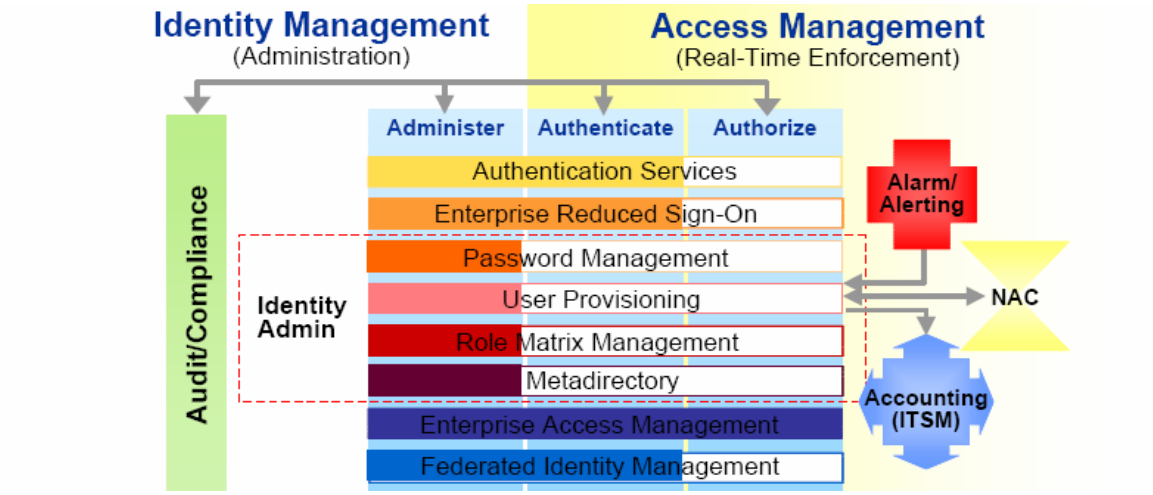

Figure 1. Gartner's breakdown of IAM market

The trouble with drawing a map is that it implies that claiming a position in every location on the map is a good thing and that everything off the map is irrelevant. As a result many of the household names in the IT industry have dashed to put a check mark in every box and to build out a 'full suite' offering. Since most have chosen the quickest path i.e. acquisition, most are still dealing with the challenge of making the various components sing in unison. It all sounds promising but in the end it may just prove to be another case of gratuitous bundling aimed at locking customers into closed implementations.

Discounting the hype around federation, the IAM market is driven by the prospect of saving money, primarily by automating what would otherwise be manual tasks such as setting up the various user accounts for new employees and getting users back to work when they forget their passwords. More recently, IAM has caught the attention of those worried about compliance, albeit largely from the point of view of generating audit trails and closing security loopholes – most importantly de-provisioning accounts for ex-employees. These are all good things but are hardly central to the vision of de-perimeterization.

In fact, the majority of the IAM market seems firmly rooted in the traditional perimeter world focused on administrative functions associated with weak password based authentication of users. Products place relatively little emphasis on security and are deployed 'safely' behind the perimeter. In some deployments security does come into play in the context of strong authentication i.e. the use of legacy tokens or smart cards, but here the emphasis is placed on authenticating select groups of remote users through the perimeter or providing physical building access to secure facilities. Rarely do these deployments extend to strengthening the collaboration and commercial activities that take place once initial access has been granted. Generally speaking, the interaction between the core IAM market and the more niche, strong authentication players has been relatively weak.

The current IAM market fails on a number of fronts to address the needs of a de-perimeterized world. It is focused on administrative tasks rather than on determining trust on which further decisions, potentially frequent decisions, can be made. It takes little or no account of context i.e. where a user is and from what machine he or she is working. Finally, it fails to recognize that the credentials within any secure system need to be better protected than the data that they themselves provide access to.

2 IAM in a de-perimeterized world

The Jericho Visioning White Paper highlights the serious issues involved in consolidating multiple IAM systems to achieve common security profiles, shared interpretations of trust and ultimately the federation of identity. However these are longer term issues for most organizations. In the short term most organizations will realize more immediate value by focusing on the needs of closed communities wishing to operate over open systems.

Organizations wishing to provide users that connect over open networks with access rights to 'behind the perimeter' systems and applications typically do so by one of two methods:

- By controlling and limiting the information flow through web portals
- By limiting the number of remote parties to those with trusted, contractually bound, relationships as is typically the case for example in most web services deployments

In situations where remote users require largely unfettered access to information the only approach is to use a higher grade of security credential, such as a hardware token. In a de-perimeterized world everyone is a remote user and in the vast majority of cases none of these approaches are either infeasible or cost prohibitive. It is necessary to consider moving to a new class of credential for everyone.

Nobody argues that the days of the password being used as a primary credential are numbered. In fact, Gartner predicts that "By 2007, 80 percent of organizations will reach the 'password breaking point' and will need to strengthen user authentication", the question is of course, what will replace passwords?

By definition the perimeterized model results in some users being inside and others outside. This results in different sets of rules, in turn creating the opportunity and even necessity for multiple and differing requirements for credentials within the same organization. In the absence of a perimeter the situation is different - everyone is equal. The emphasis of authentication shifts from validating users one by one as they cross the border and results in the widespread dependence on the use of mutual authentication. In a world where everyone can authenticate everyone else there needs to be a unified approach to credentialing otherwise costs will spiral out of control.

Cryptography has emerged as the most practical way to provide this unified approach. Cryptography has a long and proven history. Before the age of mass digital communications, cryptography was primarily used by the military for keeping valuable

information secret in hostile or open environments. In more recent times cryptography has been used in a broadly similar way by other organizations protecting intellectual property, customer information and transactional networks, particularly in situations where it crosses the Internet. Cryptographic credentials provide the sort of flexibility and security that is particularly well-suited to today's boundary-less information flows.

In understanding the management implications of the widespread use of cryptographic credentials (keys) and therefore their place in the IAM market it is important not to think of them as simply extra long passwords. The dynamics that describe their usage in a de-perimeterized world will also change compared to passwords.

2.1 Strong credentials must go beyond users

The fact that credentials will no longer be limited to gaining access to systems and applications but will be used to provide access to data, for example the ability to 'unlock' encrypted content, places greater emphasis on the user's machine. Traditionally the user's machine acted simply as the conduit through which the user's password was presented. With the widespread use of encryption for confidentiality the user's machine becomes the means to unlock the data. Therefore the degree to which the machine can be trusted to protect that content once unlocked, or at least destroy it after use, becomes an important security metric. Determining the trustworthiness of a machine is a big topic and beyond the scope of this paper but clearly the ability to validate the identity of the machine is an important early step.

Furthermore, the use of credentials will go beyond user level interaction and encompass automated systems and applications. The increasingly widespread use of SSL to protect connections between application servers and databases for example, illustrates that cryptographic credentials are already becoming pervasive. The opportunity now is to attempt to span all entities that are required to be authenticated with a unified and secure cryptographic approach.

2.2 Credentials with context

As a credential the password can deliver very little, in most cases simply resolving the question "can I cross the perimeter?" In a de-perimeterized world the controls are of a much finer grain. For example, as users and machines roam, authorization decisions will change and privileges will be adjusted, users on machines issued and controlled by the organization may very well have different rights than the same users in an Internet café or using their own PDA in a Wi-Fi hot spot. Multiple credentials will combine to create 'molecular' credentials that support more complex requests – "can this user access that data on this machine using this program, here and now?" With multiple credentials being required to enable individual transactions and collaborative tasks the number of credentials in play explodes.

2.3 Credentials for groups

Collaboration is all about groups and in a world where the subject of that collaboration is protected using encryption the context and security of the relationships and binding between the users and credentials within the group becomes as important as that of

individual user and credentials. At a more mundane level, a similar situation arises when groups of computer or programs need to access shared but protected content, for example multiple servers running different instances of the same program for resiliency or load balancing.

2.4 Rights and revocation

With so many credentials in play and with different credentials being combined depending on circumstance, the number and scope of those individuals and systems that grant rights will also grow rapidly – IT staff issuing new machines, automated patch management systems uploading security patches, business managers modifying sign-off levels and HR updating personal data. At the other end of the spectrum, the ability to revoke credentials will be required at broader level. Revocation is already a painful issue in large public key infrastructure (PKI) deployments and will need to be finally addressed. Maintaining 'black lists' of credentials doesn't scale as the number of revoked credentials can easily exceed the number that are active by an order of magnitude. With an increase in fine-grained content protection organization will need to focus on the revocation of content related access rights as well as authentication credentials, particularly as an increasing proportion of those credentials are short-lived, lasting for minutes rather than years.

2.5 Credentials need a home

Passwords tend to live in people's heads, at least they're supposed to. Sadly many get written down or consciously remembered by Windows. However, stronger credentials need a place to live where they can be protected. Once strong authentication becomes more universal and the use of stronger IT credentials moves beyond a limited subset of users a single token form factor is unlikely to be appropriate for every user. Start factoring in strong authentication for devices and the increasing availability of embedded 'tokens' such as Trusted Platform Modules[2] (TPMs) within personal computers and other devices and it is clear that the diversity of tokens in use within the organization will increase. Moreover, the mix will change over time as costs come down and functionality goes up. Therefore, our choices for credentials will need to become token independent.

3 Credential management and IAM

From this analysis it is clear that wholesale transition to stronger authentication, particularly in a way that yields the contextual detail and transaction level dynamic that the de-perimeterized requires will be no simple task. The various components of the existing IAM market address many of the challenges associated with credential usage such as account provisioning, reduced sign-on and user oriented self services but the emerging issue of credential management is currently very poorly served.

Existing approaches to strong authentication involving portable hardware tokens have a very narrow focus and, in almost all cases use local programming to configure the token

[2] A dedicated, typically hardware based security microcontroller used to protect and manage cryptographic keys and certificates. Specified by the Trusted Computing Group (www.trustedcomputinggroup.org) and implemented by leading chip and PC vendors

and the postal system to distribute the credential to the user. Obviously this doesn't map to the new world. PKI comes somewhat closer to our requirements but still falls short, focusing on certificate management rather than private key management, which of course is the actual credential. By typically relying on the end-point to generate keys PKI suffers from weak policy controls, limited audit trails, inconsistent, or completely absent, approaches to back-up and recovery and no ability to support symmetric key distribution schemes necessary for sharing secure content.

Many, if not most, of the product categories within the existing IAM market have evolved out of a need to automate manual processes (user/account provisioning, password reset) or to reduce system complexity in the eyes of the user (SSO, password synchronization, federation). As we move towards the de-perimeterized world there will need to be a step change in raw IAM security and uniformity of policy enforcement. These factors represent precursors to IAM moving from being a series of relatively infrequent administrative tasks to a position where it takes center stage as the primary trust and rights management systems across the distributed enterprise. This paper proposes that a new function, Credential Management, must be specified and implemented as a new component within the IAM set. Figure 2, below illustrates this addition to the IAM market by modifying the Gartner market breakdown diagram.

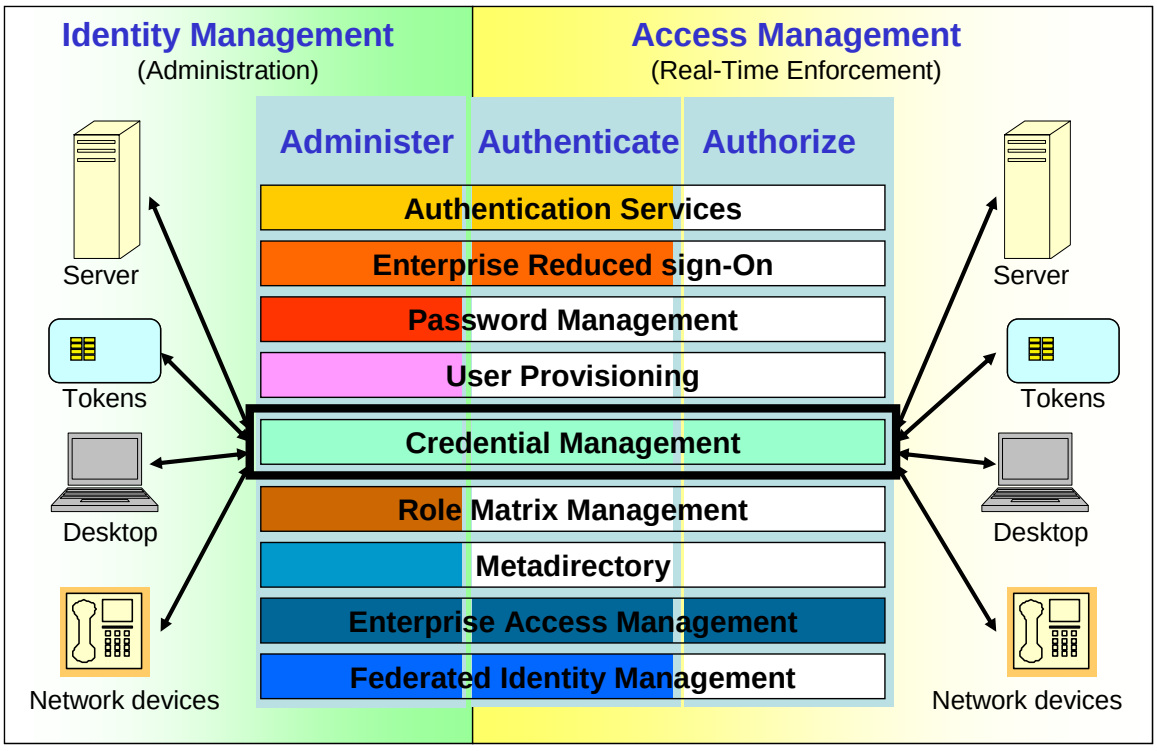

Figure 2. Modified Gartner IAM market breakdown with credential management

4 Centralized Credential Management

Centralized credential management has the potential to tackle many of the challenges associated with de-perimeterization. This paper envisages a core, centralized credential factory and repository supplying strong, cryptographic credentials, potentially on demand, via a resilient decentralized credential distribution network to a wide variety of clients and end-points represented by users, tokens, devices and applications. The system would be a definitive and secure credential system reliant on one or more authoritative sources of identity information and closely coupled with user/account provisioning work flow systems through which security profiles and templates could be created and applied. Credentials could be made mobile and able to roam between devices and would be used in conjunction with existing authentication and authorization systems, many of which already support strong cryptographic credentials.

4.1 Scale and scope

A Credential Management System of this nature clearly places a significant emphasis on scalability. As these systems become more mature they must be capable of managing and distributing volumes of credentials significantly in excess of the volumes of users managed under current IAM systems. This is for three reasons;

- The scope of credentials will apply to devices as well as users; IAM typically focuses only on users.
- Each user or device end-point is likely to have multiple credentials relating to both authentication and to access encrypted content
- Many of these credentials may be the short lived or only made available to the user for a short period of time meaning that the number of credentials actually delivered over a given period of time will be higher that traditional IAM systems where users are provisioned relatively infrequently i.e. as they join the organization or change roles.

In addition to the sheer volume of credentials, the diversity of end-points to which they are delivered will further impact the complexity of the Credential Management System. Initially the system may focus on passwords and software-based client certificates with associated software keys, distributed to modified CSPs and crypto libraries. But over time, convergence will occur with other token provisioning systems and the Credential Management System will be required to program portable hardware tokens (cards, USB etc.) as well as remotely distributing keys and other credentials to these tokens once in the field. Moreover, as trusted computers with embedded TPMs become deployed and similar approaches to internal security[3] for networking devices such as VoIP telephones and network switches reach the market, a new class of embedded token will emerge.

[3] For example Broadcom have announced the BroadSafe program (http://www.broadcom.com/press/release.php?id=525430) to embed hardware security components in a variety of networking chips and ARM have announced an embedded security technology known as TrustZone (http://www.arm.com/products/CPUs/arch-trustzone.html) for embedded microcontroller cores

Such challenges will not emerge on day one but, given the goal of achieving a unified approach to strong authentication and credentialing, the Credential Management System must have the potential to become one of the most scalable systems within the enterprise. It must also have the potential to be one of the most flexible systems. As new token types come and go and user populations migrate from today's relatively untrustworthy platforms to the next generation of hardened devices, security policy management will become increasingly complex.

4.2 Availability and performance

For many, the first steps towards de-perimeterization will involve ensuring that every connection is protected by transport layer security as a minimum. No longer restricted to Internet connections, protocols such as SSL will become widely used across the extended enterprise. In most cases these connections will rely on mutual authentication using client certificates. For security reasons these certificates will be renewed periodically. Given the threat of key-finding attacks, unless dedicated cryptographic hardware is deployed, this rollover period should be kept relatively short. It will be important that these credentials can be renewed or 'rolled' without user interaction and that the process should be performed as a background task, to avoid situations where connection requests are rejected due to invalid credentials – driving up support costs. The Credential Management System will therefore require a high degree of automation, tracking certificate lifetimes and renewing certificates and keys according to the prevailing policy.

In situations where organization have gone one step further and are encrypting individual data objects such as documents, files or database fields, credential availability becomes even more critical. In this case decryption keys will be made available to the client or end-point on request. Such delivery could potentially be required on a per action basis. Failure to provide these credentials, on-demand, will halt whatever commerce or collaborative process is taking place. Furthermore, in situations where data is stored in encrypted format there will be an additional requirement to recover such data at some time in the future. Potentially this could be many years in the future. Therefore the Credential Management System must have the ability to escrow decryption keys and to enforce a suitable recovery policy, for example through a user self-service or through specific recovery terminals under the direct control of security personnel. It is likely that the Credential Management System will feature both centralized credential management for tasks such as escrow but decentralized credential distribution to ensure resiliency and load balancing across geographies.

Couple these goals with the need to enable credentials to roam as users move within the organization and as devices and equipment using embedded credentials are redeployed within the network and it becomes even clearer that traditional IAM deployment architectures will prove inadequate.

4.3 Integration and enrollment

It is important that the Credential Management System is considered as an integral part of the overall IAM infrastructure. The integration between the Credential Management System and other IAM components such as user provisioning, certificate authorities,

directories and rights management systems is what gives the credentials value and context. For example, one the most important and resource intensive tasks in identity manage is the initial enrollment of individuals into the organization. This process requires the gathering of personal information and the definition of a policy template to reflect the rights associated with that user. Often this enrollment process will also require the involvement and approval of line managers or other operational personnel such as HR. Today the primary output of this process is the creation and provisioning of user accounts in the appropriate set of enterprise applications.

De-perimeterization will almost certainly involve a greater emphasis on enrollment as it represents the initial phase of establishing trust between users and devices. It will be critical that currently employed enrollment processes aren't simply reapplied to new enrollment tasks such as the enrollment of device level identities. For example, it is possible to envisage the enrollment of devices as an automated process with devices being supplied pre-warranted by manufactures, containing sufficient embedded credentials to enable the device to be simply plugged into the network and validated without user intervention. In situations where this is not possible it will be vitally important to establish intuitive user oriented processes to complete enrollment. Over recent years IAM systems have focused heavily on providing user oriented self-services as a way of reducing the cost of running user help desks and going forward these capabilities could be expanded to support these aspects of the Credential Management System.

One further important point of integration will be with existing or legacy credential systems. Despite the limitations of passwords, we must recognize that they will take many years to completely disappear as a primary credential (although their use as a secondary credential will clearly continue). With the goal of creating a unified credential provisioning scheme the opportunity will exist to use the Credential Management System as a proxy for converting passwords and other challenge responses based on legacy tokens into standardized cryptographic credentials. An example of such a scenario has been presented by Galwas and Peck in a separate Jericho Challenge submission[4]. Integration with existing and future certificate authorities and trusted third parties will also be an important capability. In principle the Credential Management System is complementary to a traditional certificate authority. It can help to resolve many of the traditional deployment challenges associated with PKI by addressing the specific issues of key generation, distribution and ultimately private key revocation.

4.4 Security and compliance

The Credential Management System is literally managing and protecting the keys to the kingdom, keys and credentials that are at least as valuable as the data and systems that they protect and control access to. Therefore, architectural and operational security must have a much clearer focus that is typical in current IAM deployments. Such security goes beyond purely maintaining the secrecy of cryptographic keys, some of which may need protection for years or even decades. There is also a need to ensure the integrity of the

[4] See 'A Reference Architecture for the Jericho World' - P.A. Galwas and A. Peck, nCipher Corporation – May 2005

association between different credentials, particularly in the context of group credentials and the binding of those credentials to specific rights, for example keys that may be used for decryption but not encryption.

Managing such a diversity of security policies across large volumes of users and devices is asking for anarchy without a centralized approach to policy management. It will be critical that this core component of the Credential Management System is secure, protecting keys, policies and the enforcement of those policies in a tamper-resistant environment. Given the potentially huge quantities of records to maintain it will be necessary to employ proven and scalable database technology. However, it will not be sufficient to rely on the out-of-the-box security features of the chosen database. All credentials should be encrypted within the database and all associated information should be digitally signed for integrity.

Management consoles driving the central policy engine could however be decentralized as it is likely that multiple functional groups will require access to the Credential Management System. Given the importance of establishing separation of duty and enabling dual control for administrative functions, the use of more than one console avoids the need to bring multiple individuals to the same physical location. Communications between the consoles and policy engine should be encrypted and signed.

Key and credential delivery to the end-point, potentially via one or more intermediate distribution points (for resiliency), should be via a secure protocol that can ensure confidentiality. This protocol should be granular enough to support the delivery of only the minimum required credential information and yet be flexible enough to deliver a wider set of associated credential information and other security objects if necessary. The protocol should be as lightweight as possible since many of the end-points receiving credentials may have severely limited internal processing capability. Where possible the end-point should be capable of enforcing key lifecycle and expiration periods, destroying keys as required. Ideally the end-points would be hardware protected and subjected to the same security validations as the server components within the overall system. For example, both the Trusted Computing Group and other vendor specific programs, such as Broadcom's BroadSafe, have identified FIPS[5] as an appropriate security validation.

As described earlier, one of the current market drivers within the IAM market is regulatory compliance and clearly a Credential Management System would fall under this spotlight. It will be important that all key operations – creation, delivery, recovery and destruction should be recorded in a signed and time stamped audit log for non-repudiation. These logs could then be consolidated with other auditing functions within related IAM systems to establish compliance with either internal or external requirements.

[5] See Federal Information Processing Standard: 140-2 (http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf)

5 Recommendations for the Jericho Forum

This paper has described the general requirements for and characteristics of the Credential Management System that would operate in conjunction with existing and potentially expanded components of a broader IAM infrastructure. As described earlier, there is little market evidence that the mainstream IAM market is focused on resolving these issues in a unified way. There are some activities within various standards or pseudo-standards arenas that are tackling some of these issues but most have a relatively narrow scope or are vendor driven in order to serve a particular commercial purpose. Examples of these include OATH[6], OTPS[7], SACRED[8] and TCG.

In order to make progress towards the vision to de-perimeterization it is recommended that the Jericho Forum consider the following issues.

- Suggest a generic roadmap for adoption of IAM
- Definition and standards recommendations for interoperable hardware tokens
- Recommendations and framework for device level trustworthiness – reference to IWG/TNC workgroup of TCG would be useful
- Recommendations on minimum subset of standards from PKI standard-suite
- Recommendations on deployment of mutual-authenticated SSL
- Assess group-based trust and operational relationships for business scenarios and requirements

[6] Open AuTHentication, sponsored by VeriSign – see http://www.openauthentication.org/
[7] One time passwords provisioning standard, RSA Security – see http://www.rsasecurity.com/press_release.asp?doc_id=5523&id=1034
[8] IETF RFC – see http://www.ietf.org/html.charters/sacred-charter.html