PvIB Platform voor InformatieBeveiliging

PvIB Expert Letter – March 2007

ISSN 1872-4884, Volume 3 – No. 1

March 2007

Single Sign-On for Organisations

Ernst Lopes Cardozo

Egbert Dijkgraaf

Lex Dunn

Rob Greuter

Edwin Haaring

Ernst Mellink

Eddie Michiels

Peter Rietveld

Rob van der Staaij

Gerard Zwiers

Single Sign-On (SSO) for Organisations – also called Enterprise Single Sign-On (ESSO) – is often presented as an actual product or system with specific features such as the central registration of identities and their attributes. However, SSO is actually comprised of a wide range of mechanisms that have little more in common than the fact that users usually do not have to login again when they switch between applications. In this document we shall therefore consider SSO to be a feature of strongly divergent architectures. This has consequences for the arguments that affect the business case.

We have limited our discussion to the enterprise environment, where users have a formal relationship with the organisation, such as employees, students and patients. In this type of environment it is both desirable and acceptable for users to have a single unique digital identity for all the applications. Which type of SSO is suitable for your organisation, which arguments play a role in your business case and which issues will have to be addressed during implementation?

Page

2 INTRODUCTION
- Login codes and passwords

4 DEFINITIONS AND SCOPE
- What are SSO, RESO, ESSO and WebSSO?

6 ARCHITECTURES
- Five architectures

11 ELEMENTS OF A BUSINESS CASE
- Alternatives to SSO

13 IMPLEMENTATION
- Processes and project

15 IMPORTANT ISSUES
- Source systems and standards

16 CONCLUSIONS
- Processes and project

http://www.pvib.nl/

✉ expertbrief@gvib.nl

cc SOME RIGHTS RESERVED

http://www.ibpedia.nl/

INTRODUCTION

In the beginning it was all quite simple. The first computer was set up in the bookkeeper's office and almost nobody knew how to use it. Security was limited to simple physical security. A decade later and we had the mainframe with terminals and a large number of applications. Users had to login, identify themselves using their name or employee number (login code) and authenticate this with a password. During their session, users could switch from one application to the next without having to login again. The mainframe held all the user login codes and passwords in a single file. Authorisation attributes were soon added, which specified the applications each user could access, so everything was nice and organised.

Then came the introduction of minicomputers, followed by PCs and networks. Users first had to login to their PC or network and then login again to the various applications running on the minicomputers or main frames, often using different login codes and associated passwords for each application. There were, and still are, organisations in which users require dozens of login codes and passwords. Each system has its own rules for login codes and passwords. Some passwords have to be changed every two months, and others every three months. No wonder these users keep a list close at hand that contains their login codes and passwords. These types of systems are especially hard on highly mobile people, such as employees in hospitals and at local authority counters. They are constantly moving around during the day, and therefore changing terminals, and are obliged by law to use individual identification. All these organisations are therefore looking for a system that allows their employees to login just once, or at least to minimise the number of logins.

The different login codes and passwords result in several major problems:

- The repeated logins are time consuming, irritating and demotivating for the users.
- They produce extra work for the help desk as forgotten passwords can account for a substantial part of the helpdesk workload.
- They are unsafe because users are unable to remember a multitude of passwords, so they write them down, with all the associated consequences.
- They are unsafe because passwords that are strong and change frequently are even more difficult to remember. Users and organisations will resist attempts to improve password quality. Advanced authentication methods, such as one-time passwords and smartcards, are even less acceptable when they have to be used dozens of times a day. They are also difficult to implement without centralising identification and authentication.

There are also certain drawbacks that are not directly related to the repeated logins, but do result from the underlying problem of uncoordinated identity registers:

- Maintaining identity registers in all the systems requires a great deal of extra manpower.
- In the case of an employee entering or leaving the organisation, the number of registers makes it difficult to provide or block access to each and every system in a timely fashion.
- As a proper overview of all the authorisations assigned to a person in the various systems is not available, it is difficult to implement and maintain the segregation of duties.
- Auditing the separate identity registers is time consuming and therefore costly. In addition, it does not provide an overall understanding of the combination of rights of each person.

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

- If a system needs to access another system's data, it is not possible for the requesting system to pass on the user's identity to the supplying system in order to ensure the user's authorisations are respected. This can easily lead to compromised access rules. In addition, the transaction log of the supplying system is unable to reflect the identity of the user that accessed this data.
- The previous points are more important in organisations that are subject to specific data security laws from the government, financial or health care sectors.
- With multiple identities per user it is difficult to implement user-specific settings, such as language, across all the applications.
- It is difficult to detect suspect login behaviour if there is no central point from which all login attempts can be monitored.

Each of these problems is an argument in favour of SSO. Despite these serious disadvantages, the old situation also has a few advantages:

- Spreading the risk:
  - If a username and password or even a whole identity register, is compromised, the damage will be limited to just a single system.
- Flexibility:
  - New systems, resulting perhaps from a merger or acquisition, can be immediately operational without having to integrate identification, authentication and authorisation into the existing environment.
  - It is easy to create different security levels for different systems.
- Availability:
  - Systems and applications are not dependent on the availability and accessibility of central facilities.

There are some additional generic arguments that can be raised against SSO:

- SSO products and their implementation require time, money and management attention.
- A single authentication system means the organisation and all its systems are dependent on a single supplier.
- The SSO project impacts the automation and the processes of the entire organisation.

Although these disadvantages generally do not outweigh the advantages of SSO, they are important points for any SSO project.

Definitions and Scope

Unfortunately, SSO-related terminology has not yet been standardised. However, the following terms are often used:

Single Sign-On (SSO)	The feature of a system or infrastructure that limits the number of logins required per user to just one per session.
Reduced Sign-On (RSO)	The feature of a system or infrastructure that significantly reduces the number of login moments during a session. The reason for an extra login can be technical (incompatibility of systems) or logical (separation of identity registers for reasons of security).

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

Enterprise Single Sign-On (ESSO) or Reduced Enterprise Sign-On (RESO)	The 'Enterprise' element can indicate that the system is particularly suitable for environments that include legacy applications (minicomputers, mainframes and terminal sessions). 'Enterprise' can also indicate that the system is primarily intended for closed environments, rather than systems intended for the general public such as on the Internet.
Web Single Sign-On (WebSSO)	Identification and authentication for systems that use web browsers and the http or https protocol, either on the company intranet or the public internet.

Scope

This Expert Letter focuses on the Single Sign-On within an organisation, where there is some degree of central control over policy, applications and the technical infrastructure. 'Some degree', because so many forces are at work inside large organisations that it is often illusory to think that an entire organisation will be able to adhere to a single set of rules. International organisations have to deal with different and often conflicting legislative and regulatory requirements. Mergers and acquisitions combine non-standard systems and applications, which makes it difficult to construct and maintain a homogeneous infrastructure. It is therefore not unusual to find a mix of Windows clients and servers, as well as web-based applications, terminal-based ERP systems and UNIX applications, together with identity registers (directories) from different suppliers such as Microsoft, Oracle and Novell within a single enterprise.

SSO is firmly related to the need for strong identification and authentication. However, our scope does not include the issues related to the identification and authentication of the general public, with its specific aspects of scale, privacy and the lack of standardisation and control over the workstations.

FEATURE

Whichever flavour you pick, it is important to know that SSO is not a 'system' but a *system feature*, despite the fact that many products are offered under the label of SSO and other similar terms. Further analysis shows us that there are several, quite different, configurations on offer that possess the properties of SSO to varying degrees. These products have very different architectures. Some shield the user from the diversity of underlying applications, allowing the user to login just once, while little has changed behind the scenes with each application still using its own identities, authentication and authorisation. Other products implement central Identity and Access Management (IAM), with SSO as one of the results. In the ARCHITECTURES chapter, we will explore these architectures, but let us first indulge in a little detour into an 'utopian' IAM architecture.

Utopia Ltd.

Utopia Ltd is a large, geographically diverse organisation, such as that formed by a cluster of merged hospitals. The law requires this organisation to strictly control access to patient data. However, there are also numerous other collections of digital data that come within the scope of a wide range of access policies, ranging from fully public to highly restricted.

Utopia maintains several identity registers (directories), which are databases filled with the digital identities of a range of people including employees, patients and trainees. To maintain the effectiveness of the Identity Management processes, registers have been set up for each

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

site and occasionally even some departments. The systems and applications that require access control have been built to accept identities and their attributes from each of the registers within the organisation. Applications grant fine-grained access to information and functions, based on attributes (roles) that have been assigned to an identity. Knowing the user enables them to store the user's identity alongside each entry in the transaction log. This even applies to applications that are not accessed directly by the user, but which provide services to other applications. The user's identity is passed with each request and is used to determine which information and functions are made available.

However, before this can be carried out, the user must first complete the identification and authentication process. His identity, consisting of a user and directory name, such as pete.smith@neurology, points to his identity record and allows the system to access his authentication attributes, such as a hashed password. These are used to validate the authentication data provided by the user. If successful, the system grants a session certificate, like a 'visa', that automatically identifies and authenticates the user (or his terminal?) to all applications.

In addition to the user's identity, the session certificate contains the user's 'roles', which are used by applications to determine the information and functions that are available to this person. The session certificate also contains an indication of the level of quality of the authentication that was used, which can consist of a simple password or a more secure means such as a smart card. This enables applications to differentiate their access control decisions, granting easy access to insensitive information but forcing strict control for critical items. In most cases a single sign-on will suffice, unless the user needs access to a function that requires a more secure form of authentication. A physician may access his schedule after authentication with a simple password, but will have to use two-factor authentication, such as a smart card plus pin code, to change the medication of one of his patients.

Identity Management

SSO is therefore a feature of Identity Management (IDM or IdM) systems. Within a single organisation it is desirable for each person to use a single identity, based on their relationship with the organisation, so that all actions in the physical as well as digital world can be traced back to the individual. This is an important difference to the public Internet, where a single person can have very different relationships with different parties, such as the bank, doctor, local government, tax department and shops. Here nobody wants to use a single identity in order to prevent, or at least hamper, parties from creating a personal profile. This is not a problem within the employer-employee relationship, although there may be exceptions, such as for the workers council.

Within the context of an organisation, Identity Management is comprised of creating identities and their attributes, maintaining this data, using the identities and auditing and reporting these processes. Identities are used to control and log access to physical and digital objects in accordance with standing policies. Traceability is an important aspect and if well advertised, is a strong deterrent against misconduct. Both updates to and views of a digital object can be logged. These access events can be stored in a separate log, but increasingly often they are recorded in the object itself.

Roles are used to avoid the need for each application to maintain a list of all the identities together with their access rights as well as the alternative but equally complicated need for the central directory to maintain all the details of access to information and functions for each application. If implemented properly, applications will map their access rights to a limited number of roles. This is called Role Based Access Control (RBAC).

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

As access management and traceability are the two main objectives of identity management, this activity is often called Identity and Access Management (IAM). This also includes the processes required to manage the authorisation tables in each of the applications, such as creating, updating, checking and reporting the entries. Additional information can be found in the Dutch-language document ‘Handreiking voor Identiteitenbeheer’, which is available at www.pvib.nl.

Federation

Federated Identity Management is generally defined as the ability of a user to use a single name and password to access applications from different organisations. From the point of view of the organisation, this can also be described as applications that provide access to users that are registered in different directories. These directories may be owned by unrelated organisations, such as partners in a logistics chain, or by subsidiaries within a holding company or units in an organisation. These situations can often result from a merger.

The principal idea is that each user is registered in a single directory and therefore only has a single digital identity. The user performs his authentication in his personal directory. Applications have to trust the different directories. In the case of RBAC, agreements have to be made regarding the use of roles, which may vary from an ‘guest’ role that provides certain access rights to applications up to almost full synchronisation of all roles in a merged organisation. In practical terms, it may be impossible to merge directories because of issues such as the cost of new identity tokens, the impact on the organisation or differences in legal conditions in different countries.

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

ARCHITECTURES

Products, applications and systems that implement varying levels of Single Sign-On can be classified into five architectures. This chapter describes each of these architectures and their main features. They can be so different that the SSO feature is their only common trait.

Five architectures

1. **Client-proxy software** on the workstation with a local identity table. The user will perform the authentication process against the local table, after which the client software will login the user automatically to the various applications using identities and passwords from the local table. Some implementations of the client-proxy software are able to perform background processing of requests for password renewal that arrive from an application.

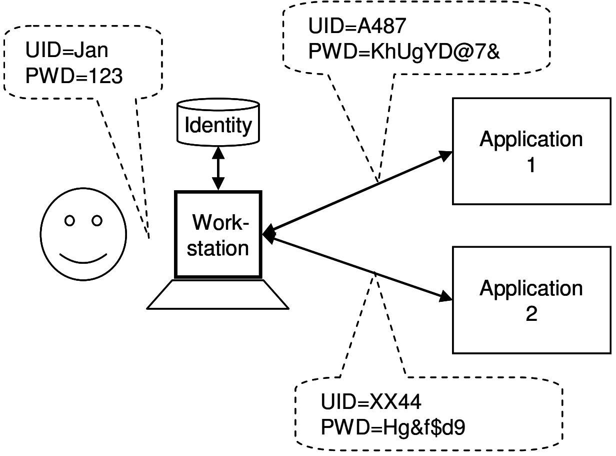

   The user's identities are only available at the user's workstations unless they are stored on a removable memory device such as a smart card. Authorisation must be handled entirely by the application servers.

2. **Synchronisation** of identities and passwords across multiple applications and systems. The synchronisation can be controlled from a central server or implemented as a distributed peer to peer system. The most basic setup only synchronises the passwords, so when the user switches from one application to the next a software component on the workstation will attempt to login automatically using his current identity. The result for the user is SSO. If all parts of an identity (name, password and roles) are synchronised in real time, we speak of an IAM system that maintains a single identity for each user. The local directories for each application can then be seen as a cache of the central directory.

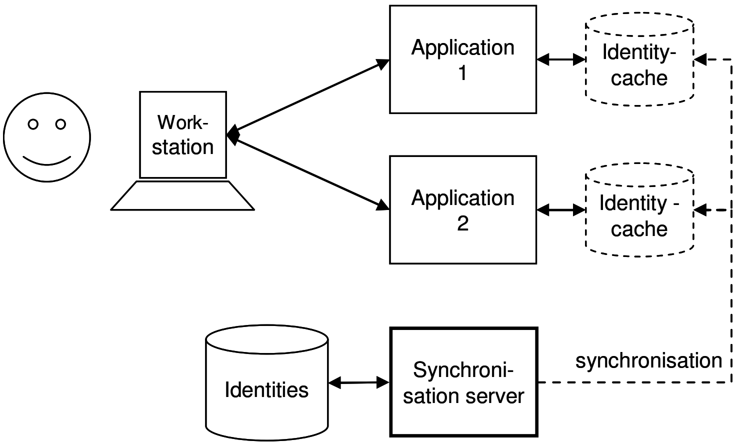

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

3. **Inline proxy**: the user logs in to the proxy. The proxy has access to a directory from which it retrieves application-specific identities and passwords to login the user to the applications.

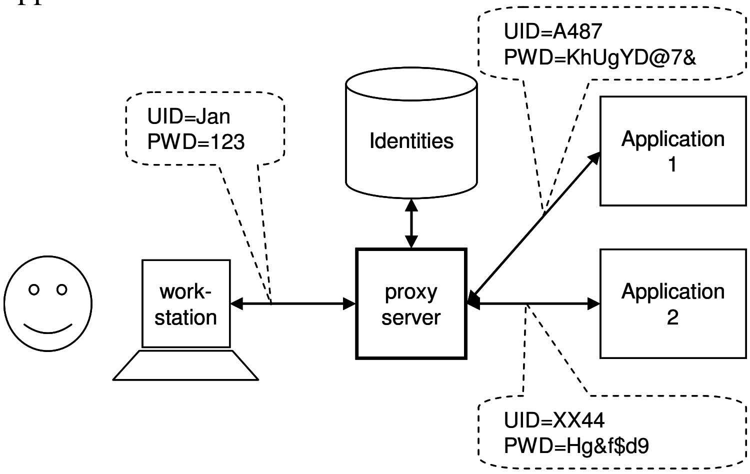

Sometimes a proxy can be used for central identity management, whereby the proxy is able to create new identities and delete old ones on the application systems. This requires a trust relationship between the proxy and the applications. It is possible to use multiple inline proxies, such as one for employees and another for clients, each having a different set of identities.

4. **Inline authentication server:** if the proxy is able to create, update and delete identities in each of the applications, it can be used for centralised identity management. This requires a trust relation between the proxy and the applications.

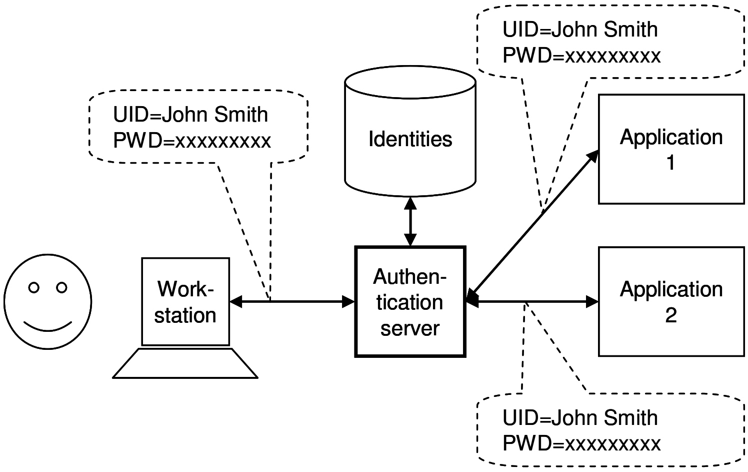

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

5. **Trusted authentication service**: a central authentication server has access to all the identities and their authentication and authorisation attributes. ① The user performs the identification and authentication procedure with the authentication server. ② The server issues a ‘ticket’ to the workstation – comparable to a visa – which is a cryptographically signed certificate from a trusted party that has validated the authentication. ③The ticket is passed to the application where it is checked to determine whether it is genuine. The ticket contains the identity of the user, which is used for logging, and his authorisation attributes, such as roles for Role Based Access Control (RBAC).

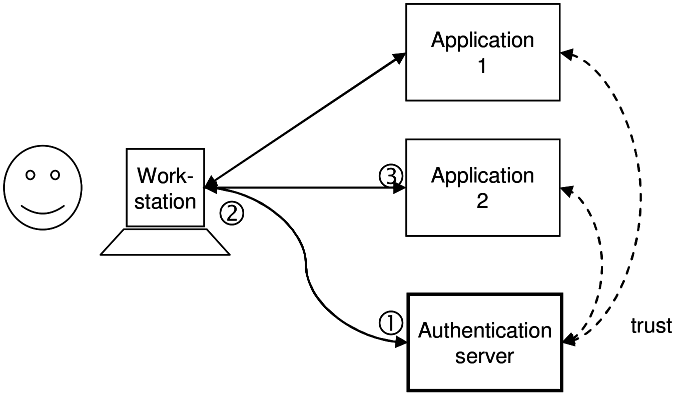

This architecture requires a certain amount of software on the workstation, which may be part of a specific client or a generic mechanism in an internet browser.

Features

We will now compare the above architectures from the perspective of SSO, in respect of the following six features:

1. Do users have a single login code for all the systems, or in other words do the identities apply to all the applications?
2. Does the architecture facilitate strong authentication methods, such as one-time passwords, smart cards and biometry?
3. Does it support Role Based Access Control (RBAC)?
4. Does it support centralised identity and access management?
5. Transparency: is it possible to prove that all identities are unique?
6. Does it support federation of identity en access management?

	Client-proxy software	**Synchronisa-tion across applications**	**Inline proxy**	**Inline authentica-tion server**	**Trusted authentica-tion service**
Single identity for all applications?	**No**	**Possibly***	**No**	**Yes**	**Yes**
Supports strong authentication?	**Yes**	**Possibly****	**Yes**	**Yes**	**Yes**
Supports Role Based Access Control?	**No**	**No**	**Yes**	**Yes**	**Yes**
Supports central identity and access management?	**No**	**Yes**	**No**	**Yes**	**Yes**
Transparency (demonstrably unique identities)?	**No**	**Yes**	**No**	**Yes**	**Yes**
Supports federated identity and access	**No**	**No**	**Yes**	**Yes**	**Yes**

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

management?					

* *Identities are only unique if synchronisation functions correctly for all attributes of the identities, not just for passwords.*

** *This poses demands on the systems and applications, which in practice are hardly ever met.*

The above table shows that only the Inline authentication server and Trusted authentication server architectures possess all the features required for proper information security now or in the future. The difference between these two is primarily the requirements they pose on the applications. The Inline authentication server communicates with the regular user and the management interfaces of the application, whereas the Trusted Authentication Server requires applications, or the platform on which they run, to use a protocol such as Kerberos.

The first three architectures (client-proxy software, password synchronisation and inline proxy) are user-friendly and support strong authentication, but lack essential elements for IAM and consequently information security in a modern application architecture. The most important limitation of the inline proxy is that in connected systems, such as the Service Oriented Architecture, there are no genuinely unique identities.

The level of ease or difficulty required to implement each of the architectures is strongly dependent on the environment, which includes the type of users, the number of clients, the number of different types of servers and applications, the geographical distribution and the organisational environment.

From a security viewpoint, the client-proxy is the weakest option as by definition the clients lack physical security and therefore should not be trusted. The other architectures can be provided with good security through acknowledged methods. One crucial step in this is the proper determination of who is authorised to create and modify identities.

Availability and reliability

Every form of centralisation increases the impact of an incident. This must be compensated. If all the systems can be accessed using a single identity, this means the impact of an identity theft will be greater. The answer to this is to reduce the risk of theft through better security for the identity data as well as stronger authentication. Fortunately, these two measures are considerably simplified by limiting the number of identities and authentications.

Storing the identity data in a single database will make the organisation vulnerable to breakdowns, data corruption and data compromise. The answer is to take measures to increase availability and security, which is again simplified because only a single database has to be protected.

A single database of identity data makes it possible to centrally monitor login attempts and therefore detect attempts at improper use. Users can learn to use their single password carefully, or to use two-factor authentication.

Naturally, authorisations must always be assigned on an individual basis, using the principle of ‘least privilege’. This is considerably easier to set up, maintain and evaluate in an environment that supports RBAC.

If the network infrastructure provides insufficient guarantees for access availability or performance, a replicated setup can provide a solution with a copy of the database at each

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

place of business. However, this requires the automatic and immediate synchronisation of each change.

The introduction of SSO can also result in the interdependence of all the applications and systems, whereby the update or replacement of a single system could affect one or more other systems.

BUSINESS CASE ELEMENTS

The above clearly demonstrates that there are two primary motives for deciding to implement a Single Sign-On system, namely to reduce costs and to improve information security through stronger identity and access management. An SSO project is often initiated because of the need to comply with new legislation. Although these motives are not incompatible, they can result in different choices and therefore different business cases.

The financial motivation

The following aspects are important primarily in respect of the motivating factor of cost reduction:

- User efficiency.
- Reduced pressure of work for the helpdesk.

These elements apply to all forms of SSO, although the benefits will differ from one organisation to another. After all, efficiency improvements do not automatically result in better operating profits.

The following aspects can result in cost reductions, but only apply to certain implementations of SSO:

- Reduced administrative burdens if the IAM is centralised.
- Prevent excessive costs for audits that are required by law.
- Decisiveness through the rapid entry and modification of identities and authorisations.
- Reduced implementation costs for new custom-made software when centralising identification, authentication and RBAC. The drawback is that existing software may need to be modified.

The primary objective of strong authentication can be efficiency. If employees frequently use a different terminal, such as care providers in a hospital or counter personnel, an almost automatic login will substantially improve productivity and quality, like a login that can respond to the proximity of a pass (token), whereby the application follows the user from terminal to terminal (fast user switching).

The information security motivation

If the main motivating factor is to improve information security, the benefits for the implementation of any form of SSO are:

- Broader acceptance of strong authentication because of fewer login moments.

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

- With strong authentication there are no passwords and hence no telephone requests to the helpdesk for a 'password reset'. Such calls are often used for 'social engineering' attacks in which a member of the helpdesk staff is tricked by a plausible story into granting access to an intruder.

The benefits can increase considerably, depending on the architecture of the chosen form of implementation:

- A unique identity for each user enables consistent access control in a software environment in which applications access each other. The following are a few examples:
  - The legal obligations for the use of personal information records by local authorities.
  - Geographic information systems that query databases *on behalf of* the user.
  - Documentation information systems that are controlled by office applications.
- The introduction of a centralised IAM means that rules can be better enforced.
- Identities, authentication and roles can be better audited.
- There is a much richer functionality for authentication and authorisation, including the ability to differentiate between authentication quality and authorisation level, depending on the application, time, user location and the type of workstation.

Alternatives to SSO

The introduction of SSO can be a costly and lengthy project, depending on the chosen objectives and solutions. The business case will not always be positive. So are there any alternative measures that can implement at least some of the benefits of SSO at lower initial and operating costs?

- Reduce the number of different platforms. SSO is often already integrated in a homogenous environment.
- 'Web-enabling' applications, which is frequently undertaken to enable remote working and to simplify software distribution, can result in a reduction in the number of login stages or at least create an environment in which SSO is easier to realise.
- Introduce a single login code per user that is valid for all the systems and applications.
- Increase the employees' awareness of the importance of information security. This improves not only the use of passwords, but also the implementation of a 'clear desk' policy and other forms of care.
- Enable users to reset their password for an application, thereby reducing the burden on the helpdesk. If the users are able to login to the network and answer a 'secret question', this can provide sufficient security to generate a new password, perhaps by sending a text message to their registered GSM phone. If this is not possible, partial automation of the workflow technology can provide some consolation.
- Consolidating applications can be an alternative, such as reducing 1100 applications to just 50.

The above measures are generally not simple or cheap to implement, but in exceptional circumstances can be a partial alternative. Most of these measures make the subsequent introduction of SSO both simpler and cheaper. However, the automated password resets will probably become superfluous if the system is subsequently equipped with SSO.

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

SELECTION OF ARCHITRECTURE AND TECHNOLOGY

So, in a practical situation, how do we decide on a particular architecture and technology? Step one is to determine the main motivating factor. Is the primary objective of SSO to increase efficiency or increase information security?

All the architectures qualify if efficiency is the sole objective. The nature and diversity of the existing systems and IAM processes will then determine the most suitable architecture. In this case, the obvious course may be to examine the technical suitability of the different solutions, but the cost and effort required to implement the necessary organisational changes may very well be greater. An inline authentication server may appear to be a good solution, until you realise this also means that all the IAM processes must be centralised. In short, proper consideration must be given to the organisational impact during the evaluation.

If the primary goal is to improve information security, only the Inline Authentication Server and the Trusted Authentication Server qualify, if necessary, in combination with client software to integrate legacy applications. The more an organisation uses Service Oriented Architecture (SOA) and Software As A Service (SAAS), the easier and more necessary it becomes to implement independent authentication servers. The starting principle of SOA is that applications operate as generic services for the entire organisation and thus there are no longer any ‘departmental applications’. If this principle is accepted, IAM centralisation should no longer be a matter for discussion. You should note that role-based authorisation could still be a departmental responsibility. For example, the council’s records office will remain responsible for the GBA data. This is the municipal personal records database that must be used by all the departments to retrieve personal information, as prescribed by law. The departments determine the roles that are assigned to their employees, and this is stored in the authentication server. The records office specifies in the GBA application which roles have access to the GBA, and to which extent.

IMPLEMENTATION

The standard principles and techniques of project management, such as the Prince2 method, apply to the introduction of SSO. Each Prince2 project starts with the business case. This has to be based on a good analysis, including the risks, impact and a realistic planning for all aspects of the project. This includes the technical implementation of the hardware and software, but will also certainly include the organisational modifications (IAM processes) as well as training for the administrators and users. By definition, SSO affects all users!

Processes

There is also a risk that SSO is treated too much as a technical project and that procedural and relational changes receive insufficient attention. There can be major organisational consequences and they can affect many people and departments, in particular if the chosen architecture includes the centralisation of identity management. A great deal of attention will have to be focused on restructuring the identity management processes. Who is allowed to enter a new identity and which checks must be performed in advance? Who determines the roles that have to be assigned to an identity? Even with a central identities file, it can still be necessary for the identity management to be performed locally to ensure it is sufficiently decisive. For example, it must be possible for new and temporary employees as well as guests and students to be entered quickly and provided with basic roles so they can get down to work. It can be necessary to provide permanent employees with delegated authority to create

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

such ad-hoc or temporary identities, which are restrained by limitations in duration, numbers and roles for the identities that are to be created.

In larger organisations, this often sparks off a discussion on the ownership and the correctness of the existing identity data. So as not to disrupt progress, each department regards its own data as sacrosanct, even though such registers are in practice seriously contaminated. Cleaning up this data can be very labour intensive and raises the question who will carry out the work and who will pay for it. To keep this process manageable, it can be worthwhile starting with an initial group involving a limited number of departments and applications and then invite or designate other groups to join. It is crucial that all identities are entered into the new, centralised database in compliance with the applicable, but perhaps new, rules. This can mean that users must report with their identification papers before they can be assigned a new identity.

Project

An architecture and business case organisation is selected depending on whether the goal focuses primarily on cost savings or also includes increased information security. The aspects of process reorganisation dealt with above play an important role in this on both the cost and benefit sides.

If the business case is clear and it validates the objectives, the first step is to design, test and implement as many as possible of the IAM-related process changes. Only then is it time to choose the product. Proper attention should be focused on whether the proposed products do actually possess the properties on which the business case is based! You should also consider whether all the applications are worth including. Applications that only have a limited lifespan or are only used by a limited number of people could perhaps better remain outside the scope of the project. If you can achieve 95% of the benefits at 70% of the cost…

As far as the rest is concerned, the introduction of SSO is a project like any other and so the same advice applies here as it would for any project: plan, structure, define milestones, celebrate successes, manage expectations, communicate, communicate and communicate. SSO affects the entire organisation and everyone will have an opinion about it. The PvIB is preparing an Expert Letter on managing aspects of information security in Prince2 projects, which will be published on www.pvib.nl in due course.

And in conclusion, ensure the benefits that have been achieved are well publicised.

SSO MANAGEMENT

After SSO has been introduced, we must safeguard the results and ensure the organisation maximises its benefits. The first concern is that eventually identity islands will still arise. With the rise of Internet, when every organisation suddenly had to get onto the Internet, projects were developed that were often entirely independent of the existing ICT organisation, and which did not relate to it in any way at all. Everything, including management, security and backup, were set up locally – or just forgotten – with all the ensuing consequences. SSO is also at risk of becoming a local project that has to realise a goal quickly and therefore proceeds outside the regular structure. The exclusive role of the central IAM must be formally documented. However, enforcement and sanctions are usually also required. In organisations that use IT service management processes, the change management process forms a good starting point for enforcement.

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

The increased inherent risks mean that careful, periodic evaluation is required in the form of an Internal Control Programme (ICP) and the formulation of a set of standards to maintain the quality of the systems and processes. In general, external auditors are enthusiastic about this approach and will perform their auditing task more quickly and cost-effectively, whether the audit is financial, legislative or regulatory. This argument applies in particular if SSO is implemented on the basis of centralised identity management.

IMPORTANT ISSUES

Reality is always more complex and recalcitrant than both theory and the product brochure. The following are a number of the potential pitfalls.

Timely source systems

During the introduction of central identity management, there is frequently the expectation that it can be fed from existing registrations, such as the personnel information system. In practice, this often proves to be very disappointing. Certainly not all ‘users’ are included in the personnel information system. The purpose of the personnel department is to process all the updates by the end of the month so the salaries can be calculated correctly. This schedule does not correspond with the need to enter new users and role changes immediately – within a few hours or faster if possible. The personnel department is often dependent on the supply of information and authorisations from the organisation, a process that is not generally easy to speed up. On the other hand, the statutory obligation to be able to produce a copy of every employee’s identification documents from their first day means these processes need to be modified. Will this obligation soon be fulfilled by the IAM process or will the existing processes in the personnel department be speeded up?

Quality of source systems

Another potential disappointment is the quality of the existing registrations, whether these are system and application directories, files in the personnel department or other sources. If different sources are combined, because one contains the names and departments but not the email addresses or room numbers, this will invariably result in problems because of differences in notation, abbreviations and file contamination. You should consider that a substantial effort will be required to initially fill the new system and to determine how discrepancies in the source systems can be corrected.

Needs differ

It would seem that every user should be grateful for SSO, but working conditions and practices can differ greatly from one person to the next. A perfect solution for one person can pose new problems for another. Is it possible to sterilise the new means of authentication? Can you use the new tokens in an explosion-risk laboratory? Does the new login procedure also work for system managers that have to resolve network problems? A test setup at a central location that is visited by as many users as possible will help to visualise these problems at an early stage.

Lock-in

As a coordinating mechanism, a system that implements SSO will by definition result in possible product and supplier dependence. You must therefore prepare for the possibility of

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

the supplier disappearing from the scene. Placing the source code in escrow is one possible measure, but if your supplier should cease operations you will probably want to change to another supplier.

Standards

Even if the supplier is still fully functional, you could be confronted with limitations in your selected product because of changes in your ICT infrastructure or organisation, such as a merger. Properly documented interfaces that meet public standards can open the way for combinations or migrations. However, just putting these standards 'on the packaging' will certainly not guarantee problem-free cooperation with other products that cite the same standards. The entire area of IAM is undergoing rapid development in 2007, which is in part driven by the increasing problems of identity theft on the internet. This has recently resulted in an agreement between the developers of OpenID and Microsoft's CardSpace. The industry appears to have reached the conclusion that there really must be a single standard. In time, this will also gain influence in the Enterprise domain. Your position regarding standardisation, and your participation in it, should play a role in your product evaluation.

CONCLUSIONS

Single Sign-on within organisations, which means not logging in any more than is strictly necessary, is an appealing characteristic for an ICT infrastructure that increases efficiency and simplifies the introduction of strong authentication. SSO can be realised through the introduction of comprehensive Identity and Access Management (IAM) systems. They provide a major contribution to information security and are crucial to organisations that must comply with specific rules, regulations and legislation, such as local authorities, health care institutions and financial service providers. However, there are also solutions that do little more than reduce the number of login moments.

SSO not only makes authentication possible, it makes it essential. The introduction of SSO is therefore the perfect time to choose for strong authentication.

Depending on the objective and the selected resources, the business case for SSO can focus solely on cost savings but also on, or even primarily on, the improvement of security.

The introduction of SSO has organisational consequences that become more important the more IAM is implemented.

The site www.ibpedia.nl (in Dutch) enables you to help further improve and share knowledge related to SSO and other information security topics. Everyone is very welcome to participate.

The expert group is very interested in the added value to you of this expert letter and welcomes your comments. You can send your reactions to expertbrief@gvib.nl. We would also appreciate an email from you if this expert letter has *not* been of value to you!

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

GLOSSARY

Glossary	
Directory	Database containing the Identities of people and objects. We have focused on the identities of people in this document.
Identity	A record in a directory. Each record contains a unique name or number, data related to Authentication, data related to access management and other data, such as department and room number.
Identification	The process used to link a user to an Identity ("I am P53871") after the user has entered a name or number.
Authentication	The process with which a user enters one or more items of authentication data, such as a password, to prove he is entitled to use the specified Identity.
Logging in	The combination of Identification and Authentication.
Authorisation	The process that provides a user with access to applications, functions and data. Roles are also frequently assigned to an Identity so that application managers can assign access to particular Roles. (John is assigned the role of Bookkeeper; the financial package is accessible to Bookkeepers, so John has access to the financial package.)
Logging	The recording of events and activities. In this context, the records will be of the activities of the users, including such information as which Identity was used and at what the time.
SSO	Single Sign-on. This is the characteristic of systems in which users only have to login once to use all their applications.
ESSO	Enterprise Single Sign-On. This is a SSO that is suitable for a large business environment.
RESO	Reduced Enterprise Sign-on, which is a less strict form of ESSO whereby users do not have to login more often than is strictly necessary to be able to use all their applications.

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.

ANNEX: APPLICABLE LICENCE

The expert letter is published under the following licence: http://creativecommons.org/licenses/by-sa/2.5/

At the time of writing, this page appears as follows:

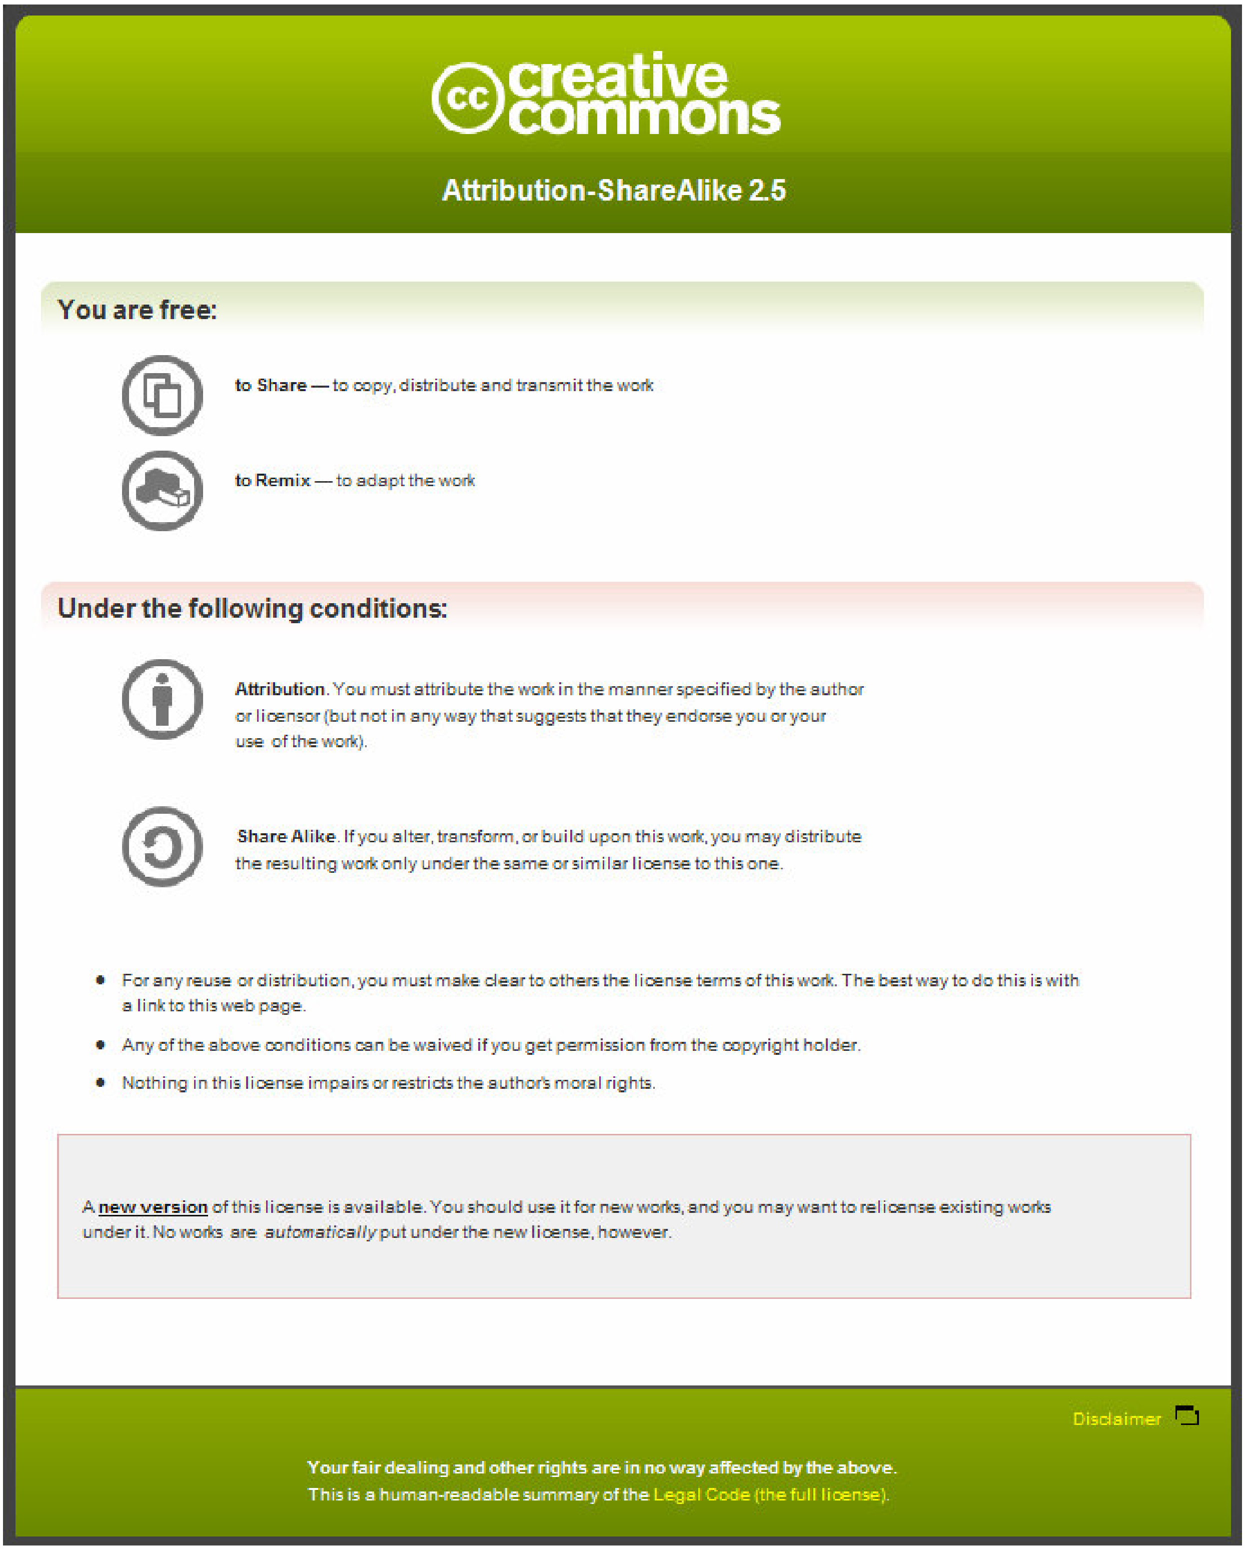

creative commons

Attribution-ShareAlike 2.5

You are free:

to Share — to copy, distribute and transmit the work

to Remix — to adapt the work

Under the following conditions:

Attribution. You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).

Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under the same or similar license to this one.

- For any reuse or distribution, you must make clear to others the license terms of this work. The best way to do this is with a link to this web page.
- Any of the above conditions can be waived if you get permission from the copyright holder.
- Nothing in this license impairs or restricts the author's moral rights.

A new version of this license is available. You should use it for new works, and you may want to relicense existing works under it. No works are *automatically* put under the new license, however.

Disclaimer

Your fair dealing and other rights are in no way affected by the above.
This is a human-readable summary of the Legal Code (the full license).

(cc) 2007, Platform for Information Security (PvIB)
This work is licensed under a Creative Commons License.

JOIN THE PVIB, FOR SAFETY AND SECURITY ...

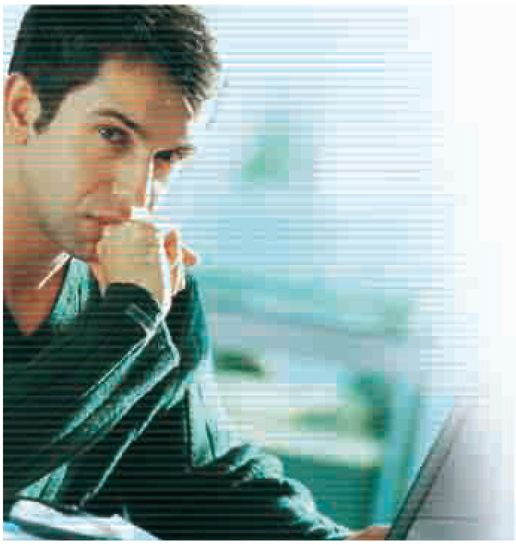

Information security has been an essential and exciting subject for many years. Almost all occupations are having to place more emphasis on the confidentiality, availability and integrity of their information. Whether you are a CISO, manager, consultant or programmer, the Platform for Information Security (PvIB) can help you with information security issues.

What is the Platform for Information Security?
The PvIB is an open, broad-based association for professionals to build a more professional approach to information security, through the exchange of ideas, information, knowledge, insights and above all, practical experience.

What are our aims?
We aim to promote the physical, systems and organisational security of data and data processing equipment against in-coming and outgoing breaches. We also promote the exchange of knowledge and experience and the networking of practitioners in the sector - through this Expert Letter, for example.

Our target group
The target group for the PvIB includes everyone involved in information security, either as a student or professionally, or who are especially interested in the field. Our rapidly growing membership covers many different disciplines: students, information architects, technicians, managers, organisational consultants, legal specialists, security officials and ICT auditors. Our members come from all kinds of educational backgrounds, companies, public authorities, organisations and suppliers.

(cc) 2007, Platform for Information (PvIB)
This work is licensed under a Creative Commons License.