

PKI applications (C2)

From IPsec to Identity management

Pascal Steichen (MSSI-uni.lu) - 25/01/2007

- 1. Implementing PKI
 - 1.1. IPsec - securing the network
 - 1.2. SSL/TLS - securing the web
 - 1.3. S/MIME - securing e-mail
 - 1.4. XML-Dsig - securing (XML) documents
 - 1.5. Securing your own applications
- 2. Authentication
 - 2.1. Authentication in PKI
 - 2.2. Schneier on Security - The Failure of Two-Factor Authentication
 - 2.3. Beyond authentication
- 3. Single Sign On - The PKI killer application ?
 - 3.1. Kerberos
 - 3.1.1. Kerberos - operation
 - 3.1.2. MS kerberos usage
 - 3.2. Other SSO implementation examples
 - 3.2.1. Windows Live ID
 - 3.2.2. WS-Trust
 - 3.2.3. Liberty Alliance
 - 3.2.4. Google
- 4. Identity management
- 5. Bibliographic references

1. Implementing PKI

PKI enables confidentiality, integrity, authenticity and non-repudiation of data exchanges, by:

- encrypting/decrypting data
- signing messages
- verification of signatures
- verification of certificate validity

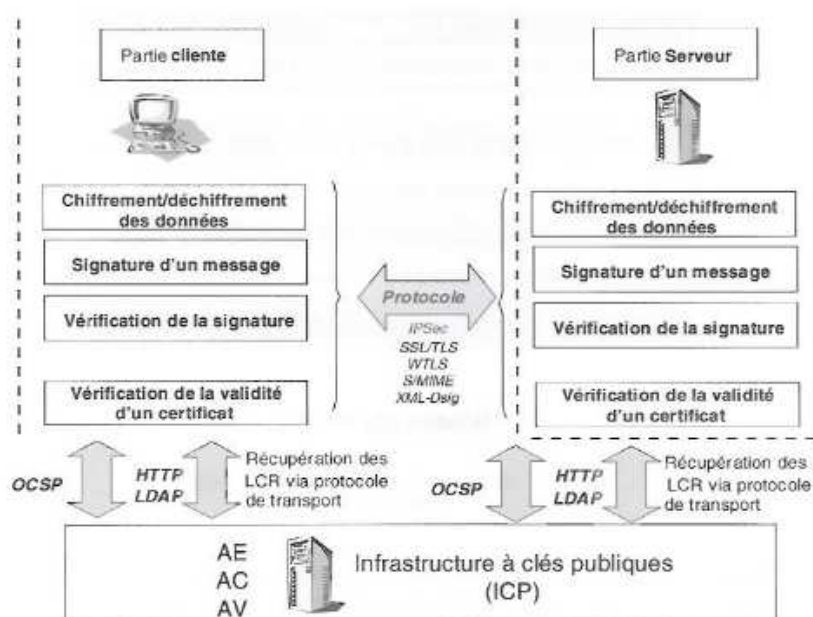


Figure 5-1. Mécanismes de sécurité et protocoles

© 2002 Editions Eyrolles, ISBN 2-212-11045-6

Implementation of security mechanisms based on X.509 certificates (and as such PKI) can be realised on different levels of an IT infrastructure (from the OSI model):

- on the network level (e.g. IPsec)
- on the transport level (e.g. SSL/TLS)
- on the the application level (e.g. S/MIME or XML-Dsig)

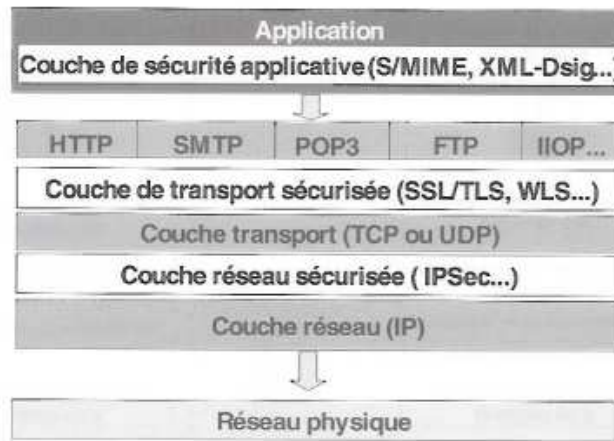


Figure 5-2. Sécurisation des couches réseaux, transport et applicatives

© 2002 Editions Eyrolles, ISBN 2-212-11045-6

As by the nature of the OSI model the different levels are independent and as such the security mechanisms can be used independently.

1.1. IPSec - securing the network

The IPsec protocol, developed by the IETF, defines a set of specifications to secure (via encapsulation) IP packets. On the network (IP) level IPsec implements a secure tunnel, also called VPN (Virtual Private Network).

To provide interoperable, high quality, cryptographically-based security IPsec is:

- designed for IPv4 and IPv6
- and offers:
 - access control,
 - integrity,
 - data origin authentication,
 - detection and rejection of replays,
 - confidentiality.

IPsec security services are provided through:

- 2 traffic protocols:
 - AH (Authentication Header)

The IP Authentication Header (AH) offers integrity and data origin authentication, with optional anti-replay features. Based on a symmetric algorithm a MAC value of the principal header fields and the data payload, is generated.

- ESP (Encapsulating Security Payload)

The Encapsulating Security Payload (ESP) protocol offers the same set of services, and also offers confidentiality. In contrast to AH does ESP only secure the data payload.

- 2 management constructs to enforce boundary security:
 - SA (Security Association)

It is necessary to manage a series of session parameters on each IPsec endpoint:

- functionalities (authentication, encryption)
- encryption modes (tunnel, transport)
- algorithms, key validity
- initialisation parameters (IV, sequence number, anti-relay, ...)

The set of these parameters defines the security association (SA), identified by a 32-bit field in the IP header.

- SPD (Security Policy Database)

The security policy database (SPD) specifies what services are to be offered to IP datagrams and in what fashion:

- if IPsec is to be applied and how (what SA, AH or ESP, etc.)
- if packets can be transferred without IPsec protection
- if they should be rejected

As such SPD defines one or more criteria, called "selectors" and analysis packets via these selectors, that can be based on IP addresses, protocol types, ports, ...

- 2 key management procedures:
 - "manual" mode

If the SA is done statically ("manual" mode), a person manually configures each system with keying material and SA management data relevant to secure communication with other systems.

- IKE (Internet Key Exchange) mode

IKE mode implies an automated SA, by providing mechanisms to negotiate algorithms and generate symmetric keys. Beforehand endpoint authentication is needed and can be done via a pre-shared key or based on a public key certificates signed by a CA. For this last option a PKI setup is of course necessary.

1.2. SSL/TLS - securing the web

SSL (Secure Socket Layer) is an open protocol securing "socket" type communication channels (e.g. TCP). Developed by Netscape, SSL is nowadays replaced by the IETF standardised protocol TLS (Transport Layer Security). Even if SSL and the later TLS are not strictly interoperable, we will use them as synonyms here after.

The protocol is composed of two layers:

- The TLS Record Protocol

At the lowest level, layered on top of some reliable transport protocol (e.g. TCP), is the TLS Record Protocol. The TLS Record Protocol provides connection security that has two basic properties:

- The connection is private.

Symmetric cryptography is used for data encryption (e.g. DES, RC4,...). The keys for this symmetric encryption are generated uniquely for each connection and are based on a secret negotiated by another protocol (such as the TLS Handshake Protocol).

- The connection is reliable.

Message transport includes a message integrity check using a keyed MAC. Secure hash functions (e.g. SHA, MD5,...) are used for MAC computations.

- The TLS Handshake Protocol.

The TLS Handshake Protocol provides connection security that has three basic properties:

- The peer's identity can be authenticated using asymmetric, or public key, cryptography (e.g. RSA).

This authentication can be made optional, but is generally required for at least one of the peers.

- The negotiation of a shared secret is secure.

The negotiated secret is unavailable to eavesdroppers, and for any authenticated connection the secret cannot be obtained, even by an attacker who can place himself in the middle of the connection.

- The negotiation is reliable.

No attacker can modify the negotiation communication without being detected by the parties to the communication.

An SSL client-server session is established via the following mechanisms:

1. The client sends session establishment data like encryption parameters, SSL version, etc. as well as a random generate code.
2. The server responds with the same type of data as well his public key certificate. Optionally it can ask for a client certificate.
3. The client identifies the server via its certificate.
4. The client then sends a "pre" secret key encrypted with the server's public key. If client auth is used, we add a signed data bloc as well as the client's certificate.
5. The server validates the client certificate if needed, then decrypts the "pre" secret key.
6. Client and server realise a series of operations to agree on the final secret key.
7. The session is then completely established.

1.3. S/MIME - securing e-mail

The S/MIME protocol (Secure Multi-purpose Internet Mail Extension), enables **digital signatures and encryption of MIME formatted messages** (e.g. e-mail). Similar to SSL, S/MIME uses a **hybrid crypto system**. The message is encrypted with a **symmetric session key**, that itself is secured by using the **destination's public key**. With multiple recipients the session key has to be encrypted several times using the different public keys.

- RFC3852 Cryptographic Message Syntax (CMS)
- RFC3851 Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification
- RFC3850 Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Certificate Handling
- RFC2631 Diffie-Hellman Key Agreement Method

1.4. XML-Dsig - securing (XML) documents

To enable secure B-to-B transactions or document exchanges over the Internet, the W3C and IETF started the XML-Signature project, defining **non-repudiation** mechanisms for signing XML (eXtensible Markup Language) documents.

The following example is a detached signature of the content of the HTML4 in XML specification:

```
<Signature Id="MyFirstSignature" xmlns="http://www.w3.org/2000/09/xmldsig#">
  <SignedInfo>
    <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xml-c14n-20010315"/>
    <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#dsa-sha1"/>
      <Reference URI="http://www.w3.org/TR/2000/REC-xhtml1-20000126/">
        <Transforms>
          <Transform Algorithm="http://www.w3.org/TR/2001/REC-xml-c14n-20010315"/>
        </Transforms>
        <DigestMethod Algorithm="http://www.w3.org/2000/09/xmldsig#sha1"/>
        <DigestValue>j6lwx3rvEP00vKtMup4NbeVu8nk=</DigestValue>
      </Reference>
    </SignedInfo>
    <SignatureValue>MC0CFFrVlTrlk=...</SignatureValue>
  <KeyInfo>
    <KeyValue>
      <DSAKeyValue>
        <P>...</P><Q>...</Q><G>...</G><Y>...</Y>
      </DSAKeyValue>
    </KeyValue>
  </KeyInfo>
</Signature>
```

Signatures are related to data objects via URIs. Within an XML document, signatures are related to local data objects via

fragment identifiers. Such local data can be included within an enveloping signature or can enclose an enveloped signature. Detached signatures are over external network resources or local data objects that reside within the same XML document as sibling elements; in this case, the signature is neither enveloping (signature is parent) nor enveloped (signature is child).

The SignedInfo element is the information that is actually signed. Core validation of SignedInfo consists of two mandatory processes: validation of the signature over SignedInfo and validation of each Reference digest within SignedInfo. Note that the algorithms used in calculating the SignatureValue are also included in the signed information while the SignatureValue element is outside SignedInfo.

1.5. Securing your own applications

Beyond using SSL/TLS, S/MIME or XML-Dsig, securing your own applications using PKI means **integration**. This can be done via 3 approaches:

- "clean mains" solution

getting a complete PKI enable system with all the needed functionalities;

- passive integration

by adding an complementary application module doing the crypto stuff and using the PKI services. In this one does not need to change the code of the application. The set-up can be done in 2 ways:

- integrating a "plug-in" or middleware
- adding a relay server ("reverse-proxy") between client and server

- active integration

by implementing the crypto and PKI stuff in the code. This implies that the CSP provides adequate APIs, 3 approaches are seen in the market these days:

- application oriented

these are development kits implementing standardised protocols like SSL/TLS, S/MIME or XML-Dsig;

- low level

these only provide the cryptographic functions needed to hash, encrypt and encode messages;

- so called "PKI" API

somehow a hybrid between the passive middleware and a real API. They simply provide the access to the PKI services, via a kind of "web service" approach.

2. Authentication

Authentication

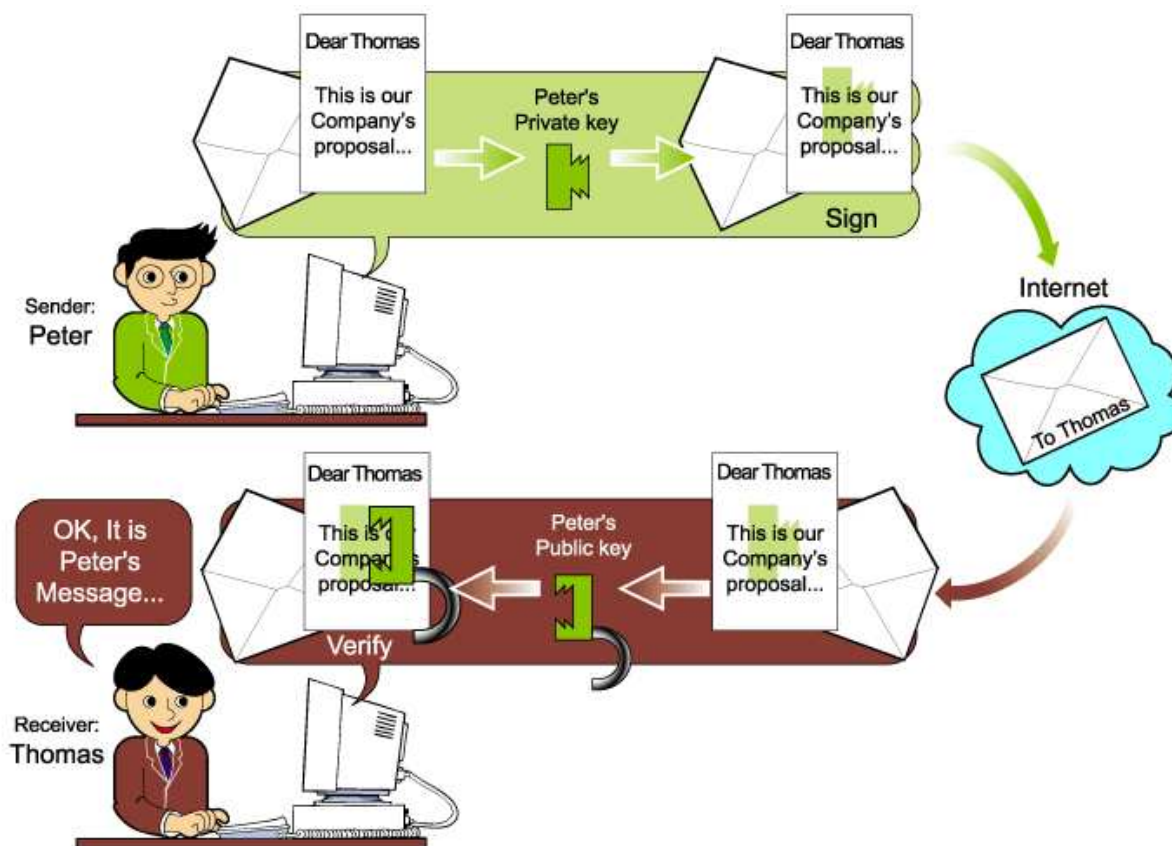
(from Greek authentikos: real or genuine; from authentes: author)

is the act of establishing or confirming something (or someone) as authentic, that is, that claims made by or about the thing are true. Authenticating an object may mean confirming its provenance, whereas authenticating a person often consists of verifying their identity. Authentication depends upon one or more authentication factors.

- The three most commonly recognized factors are:
 - 'something you know' (password, PIN, ...),
 - 'something you have' (credit card, hardware token, ...),
 - 'something you are' (fingerprint, retinal pattern or other biometric).
- Other, less common factors may include:
 - Recognition-based or cognometric authentication
 - where the user has to recognize pre-assigned secret faces
 - Cybermetric authentication
 - Only allowing access from the certain computer, which is the combination of unique hardware and (or) software installed
 - Location-based authentication
 - Only allowing a particular atm, charge, or credit card to be used at a specific merchant or at a specific bank branch, or only allowing root access from specific terminals
 - Time-based authentication
 - Only allowing access from certain accounts during normal working hour
 - Size-based authorization
 - Only allowing a specific transaction to be for a specific exact amount
 - Pre-authorized transactions
 - Where a company uploads all of the check numbers and amounts written for each check to their bank, and the bank would then reject any check not of those numbers and amounts as fraudulent

Using more than one factor is called strong authentication.

2.1. Authentication in PKI



Department of Computer Science and Information Systems, The University of Hong Kong

2.2. Schneier on Security - The Failure of Two-Factor Authentication

March 15, 2005

Two-factor authentication isn't our saviour. It won't defend against phishing. It's not going to prevent identity theft. It's not going to secure on-line accounts from fraudulent transactions. It solves the security problems we had ten years ago, not the security problems we have today.

The problem with passwords is that they're too easy to lose control of. People give them to other people. People write them down, and other people read them. People send them in e-mail, and that e-mail is intercepted. People use them to log into remote servers, and their communications are eavesdropped on. They're also easy to guess. And once any of that happens, the password no longer works as an authentication token because you can't be sure who is typing that password in.

Two-factor authentication mitigates this problem. If your password includes a number that changes every minute, or a unique reply to a random challenge, then it's harder for someone else to intercept. You can't write down the ever-changing part. An intercepted password won't be good the next time it's needed. And a two-factor password is harder to guess. Sure, someone can always give his password and token to his secretary, but no solution is foolproof.

These tokens have been around for at least two decades, but it's only recently that they have gotten mass-market attention. AOL is rolling them out. Some banks are issuing them to customers, and even more are talking about doing it. It seems that corporations are finally waking up to the fact that passwords don't provide adequate security, and are hoping that two-factor authentication will fix their problems.

Unfortunately, the nature of attacks has changed over those two decades. Back then, the threats were all passive: eavesdropping and off-line password guessing. Today, the threats are more active: phishing and Trojan horses.

Here are two new active attacks we're starting to see:

- * Man-in-the-Middle attack. An attacker puts up a fake bank website and entices user to that website. User types in his password, and the attacker in turn uses it to access the bank's real website. Done right, the user will never realize that he isn't at the bank's website. Then the attacker either disconnects the user and makes any fraudulent transactions he wants, or passes along the user's banking transactions while making his own transactions at the same time.

- * Trojan attack. Attacker gets Trojan installed on user's computer. When user logs into his bank's website, the attacker piggybacks on that session via the Trojan to make any fraudulent transaction he wants.

See how two-factor authentication doesn't solve anything? In the first case, the attacker can pass the ever-changing part of the password to the bank along with the never-changing part. And in the second case, the attacker is relying on the user to log in.

The real threat is fraud due to impersonation, and the tactics of impersonation will change in response to the defences. Two-factor authentication will force criminals to modify their tactics, that's all.

Recently I've seen examples of two-factor authentication using two different communications paths: call it "two-channel authentication." One bank sends a challenge to the user's cell phone via SMS and expects a reply via SMS. If you assume that all your customers have cell phones, then this results in a two-factor authentication process without extra hardware. And even better, the second authentication piece goes over a different communications channel than the first; eavesdropping is much, much harder.

But in this new world of active attacks, no one cares. An attacker using a man-in-the-middle attack is happy to have the user deal with the SMS portion of the log-in, since he can't do it himself. And a Trojan attacker doesn't care, because he's relying on the user to log in anyway.

Two-factor authentication is not useless. It works for local login, and it works within some corporate networks. But it won't work for remote authentication over the Internet. I predict that banks and other financial institutions will spend millions outfitting their users with two-factor authentication tokens. Early adopters of this technology may very well experience a significant drop in fraud for a while as attackers move to easier targets, but in the end there will be a negligible drop in the amount of fraud and identity theft.

© 2005 Bruce Schneier

2.3. Beyond authentication

The 4-As:

- Authentication

Authentication refers to the confirmation that a user who is requesting services is a valid user of the network services requested. Authentication is accomplished via the presentation of an identity and credentials. Examples of types of credentials are passwords, one-time tokens, digital certificates, and phone numbers (calling/called).

- Authorization

Authorization refers to the granting of specific types of service (including "no service") to a user, based on their authentication, what services they are requesting, and the current system state. Authorization may be based on restrictions, for example time-of-day restrictions, or physical location restrictions, or restrictions against multiple logins by the same user. Authorization determines the nature of the service which is granted to a user. Examples of types of service include, but are not limited to: IP address filtering, address assignment, route assignment, QoS/differential services, bandwidth control/traffic management, compulsory tunnelling to a specific endpoint, and encryption.

- Accounting

Accounting refers to the tracking of the consumption of network resources by users. This information may be used for management, planning, billing, or other purposes. Real-time accounting refers to accounting information that is delivered concurrently with the consumption of the resources. Batch accounting refers to accounting information that is saved until it is delivered at a later time. Typical information that is gathered in accounting is the identity of the user, the nature of the service delivered, when the service began, and when it ended.

- Audit

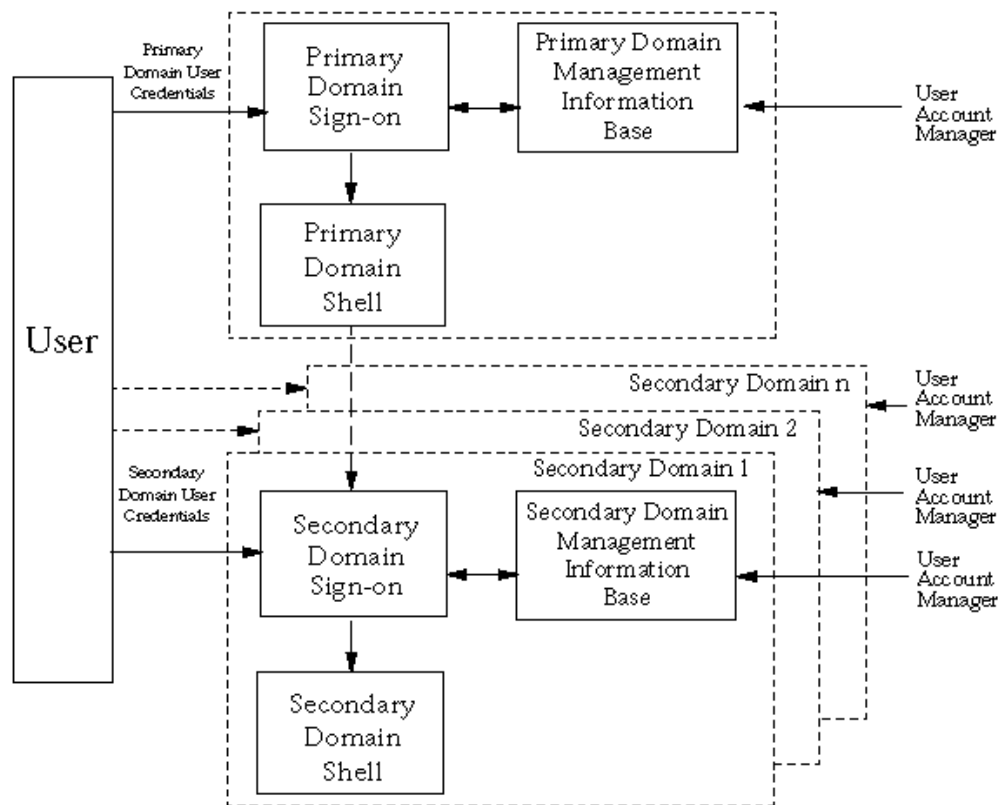
The AAA is sometimes combined with auditing and accordingly becomes AAAA

3. Single Sign On - The PKI killer application ?

A specialized form of authentication that enables a user to authenticate **once** and gain access to the **multiple** resources.

As IT systems proliferate to support business processes, users and system administrators are faced with an increasingly complicated interface to accomplish their job functions. Users typically have to sign-on to multiple systems, necessitating an equivalent number of sign-on dialogues, each of which may involve different usernames and authentication information. System administrators are faced with managing user accounts within each of the multiple systems to be accessed in a co-ordinated manner in order to maintain the integrity of security policy enforcement. This legacy approach to user sign-on to multiple systems is illustrated below:

Legacy Approach to User

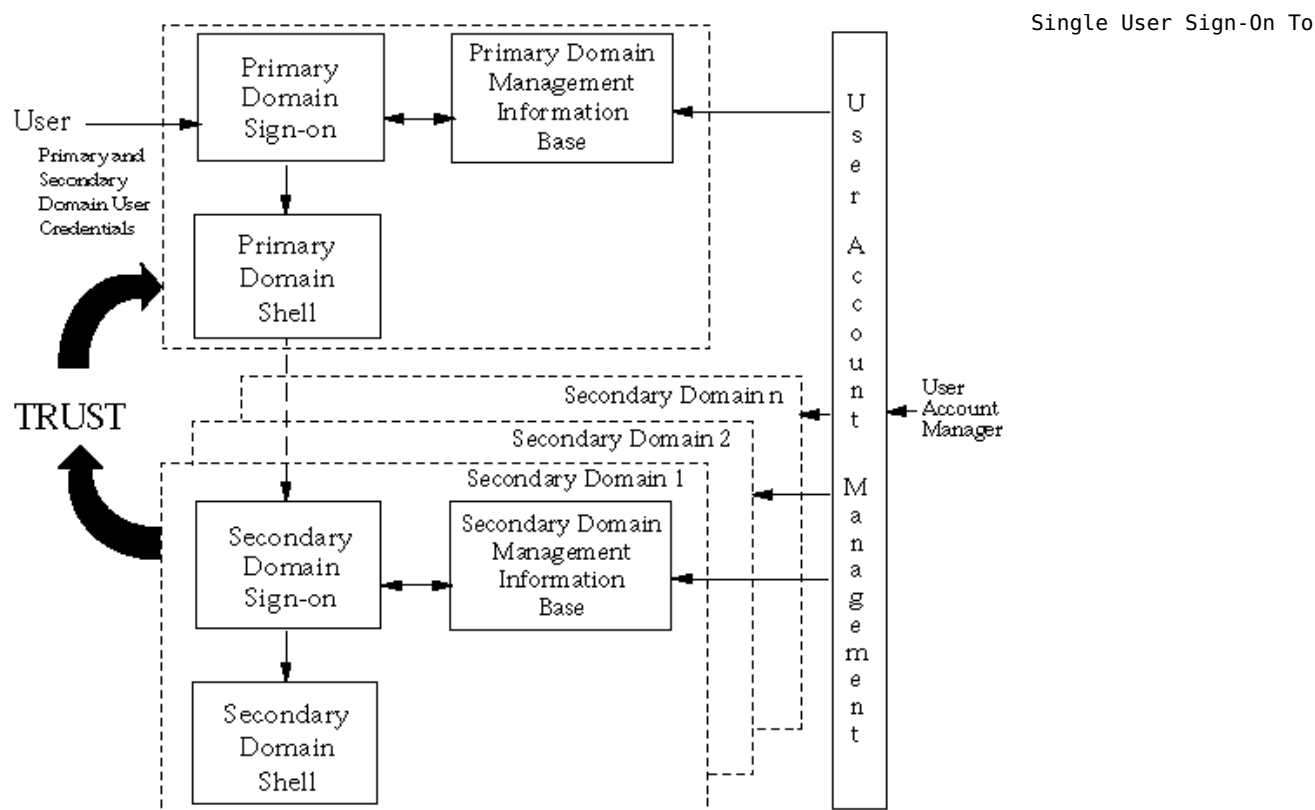


Sign-on to Multiple Systems - © 1995-2005 The OpenGroup

Historically a distributed system has been assembled from components that act as independent security domains. These components comprise individual platforms with associated operating system and applications.

These components act as independent domains in the sense that an end-user has to identify and authenticate himself independently to each of the domains with which he wishes to interact. This scenario is illustrated above. The end user interacts initially with a Primary Domain to establish a session with that primary domain. This is termed the Primary Domain Sign-On in the above diagram and requires the end user to supply a set of user credentials applicable to the primary domain, for example a username and password. The primary domain session is typically represented by an operating system session shell executed on the end user's workstation within an environment representative of the end user (e.g., process attributes, environment variables and home directory). From this primary domain session shell the user is able to invoke the services of the other domains, such as platforms or applications.

To invoke the services of a secondary domain an end user is required to perform a Secondary Domain Sign-on. This requires the end user to supply a further set of user credentials applicable to that secondary domain. An end user has to conduct a separate sign-on dialogue with each secondary domain that the end user requires to use. The secondary domain session is typically represented by an operating system shell or an application shell, again within an environment representative of the end user. From the management perspective the legacy approach requires independent management of each domain and the use of multiple user account management interfaces. Considerations of both usability and security give rise to a need to co-ordinate and where possible integrate user sign-on functions and user account management functions for the multitude of different domains now found within an enterprise. A service that provides such co-ordination and integration can provide real cost benefits to an enterprise through:



Multiple Services - © 1995-2005 The OpenGroup

- reduction in the time taken by users in sign-on operations to individual domains, including reducing the possibility of such sign-on operations failing
- improved security through the reduced need for a user to handle and remember multiple sets of authentication information,
- reduction in the time taken, and improved response, by system administrators in adding and removing users to the system or modifying their access rights,
- improved security through the enhanced ability of system administrators to maintain the integrity of user account configuration including the ability to inhibit or remove an individual user's access to all system resources in a co-ordinated and consistent manner.

Such a service has been termed Single Sign-On after the end-user perception of the impact of this service. However, both the end-user and management aspects of the service are equally important. This approach is illustrated in the diagram above. In the single sign-on approach the system is required to collect from the user as, part of the primary sign-on, all the identification and user credential information necessary to support the authentication of the user to each of the secondary domains that the user may potentially require to interact with. The information supplied by the user is then used by Single Sign-On Services within the primary domain to support the authentication of the end user to each of the secondary domains with which the user actually requests to interact.

The information supplied by the end-user as part of the Primary Domain Sign-On procedure may be used in support of secondary domain sign-on in several ways:

- **directly**, the information supplied by the user is passed to a secondary domain as part of a secondary sign-on,
- **indirectly**, the information supplied by the user is used to retrieve other user identification and user credential information stored within the single sign-on management information base. The retrieved information is then used as the basis for a secondary domain sign-on operation,
- **immediately**, to establish a session with a secondary domain as part of the initial session establishment. This implies that application clients are automatically invoked and communications established at the time of the primary sign-on operation,
- **temporarily** stored or cached and used at the time a request for the secondary domain services is made by the end-user.

From a management perspective the single sign-on model provides a single user account management interface through which all the component domains may be managed in a coordinated and synchronised manner.

Significant security aspects of the Single Sign-On model are:

- the secondary domains have to trust the primary domain to:
 - correctly assert the identity and authentication credentials of the end user,
 - protect the authentication credentials used to verify the end user identity to the secondary domain from unauthorised use,
- the authentication credentials have to be protected when transferred between the primary and secondary domains against threats arising from interception or eavesdropping leading to possible masquerade attacks.

3.1. Kerberos



MIT developed Kerberos to protect network services provided by Project Athena. The protocol was named after the Greek mythological character Kerberos (or Cerberus), known in Greek mythology as being the monstrous three-headed guard dog of Hades.

- Kerberos is a computer network authentication protocol

which allows individuals communicating over an insecure network to prove their identity to one another in a secure manner. Kerberos prevents eavesdropping or replay attacks, and ensures the integrity of the data. Its designers aimed primarily at a client-server model, and it provides mutual authentication - both the user and the server verify each other's identity.

- Version 5 is now (2005) available as RFC 4120.
- Kerberos builds on symmetric key cryptography and requires a trusted third party.

Kerberos uses as its basis the Needham-Schroeder protocol. It makes use of a trusted third party, termed:

- Key Distribution Center (KDC)

which consists of two logically separate parts:

- Authentication Server (AS)
- Ticket Granting Server (TGS)

providing 'access' to

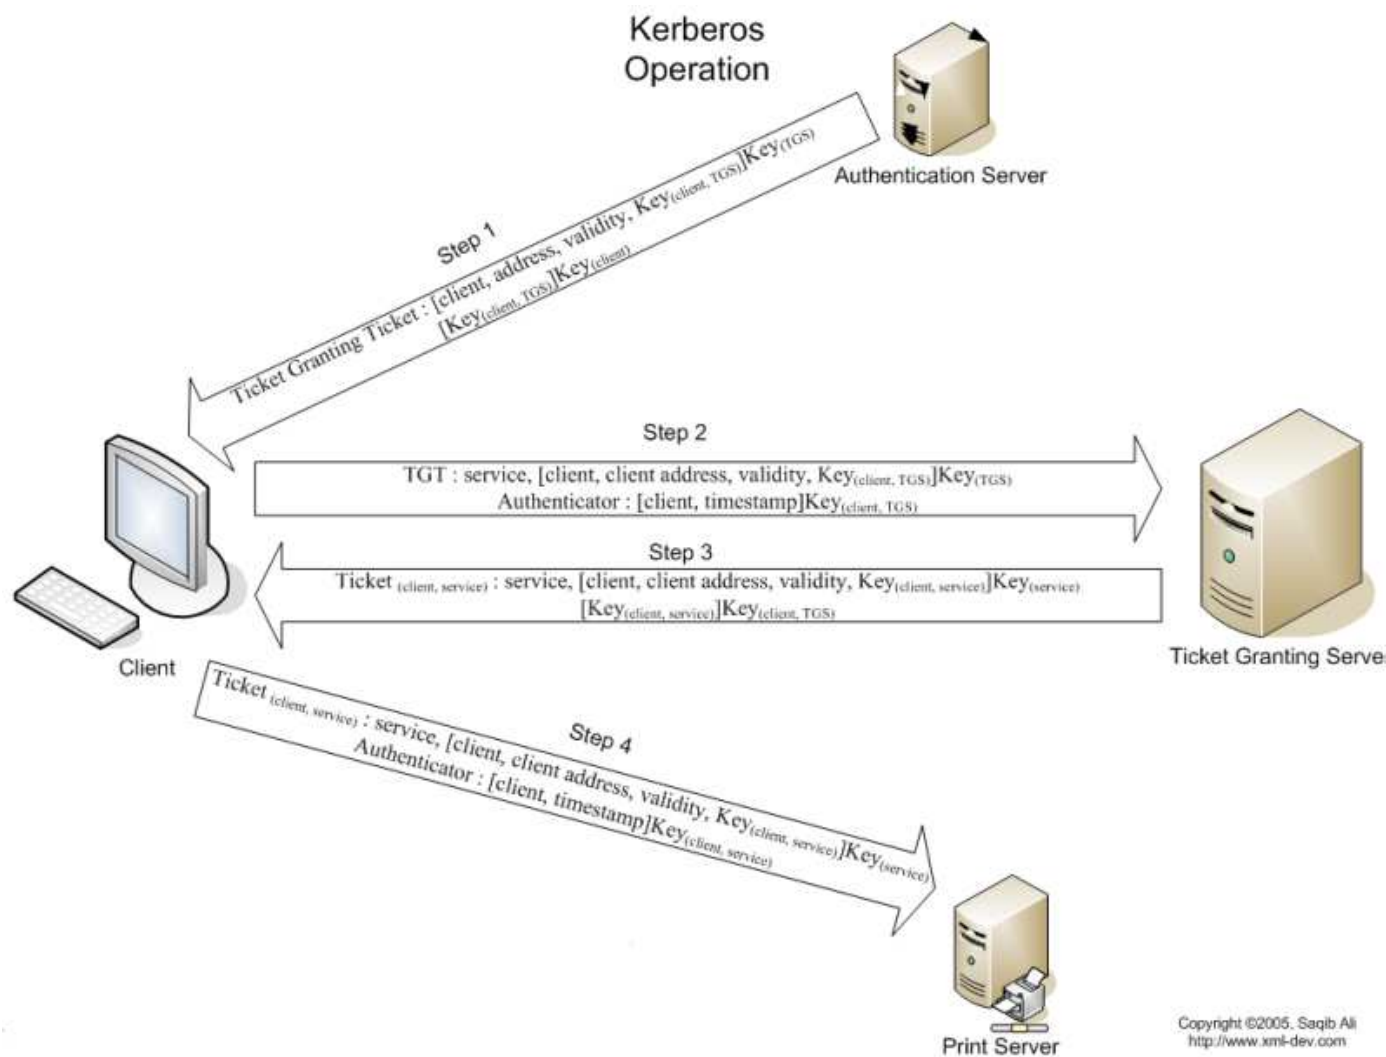
- Service Server (SS)

Kerberos works on the basis of "tickets" which serve to prove the identity of users.

The KDC maintains a database of secret keys; each entity on the network - whether a client or a server - shares a secret key known only to itself and to the KDC. Knowledge of this key serves to prove an entity's identity. For communication between two entities, the KDC generates a session key which they can use to secure their interactions.

3.1.1. Kerberos - operation

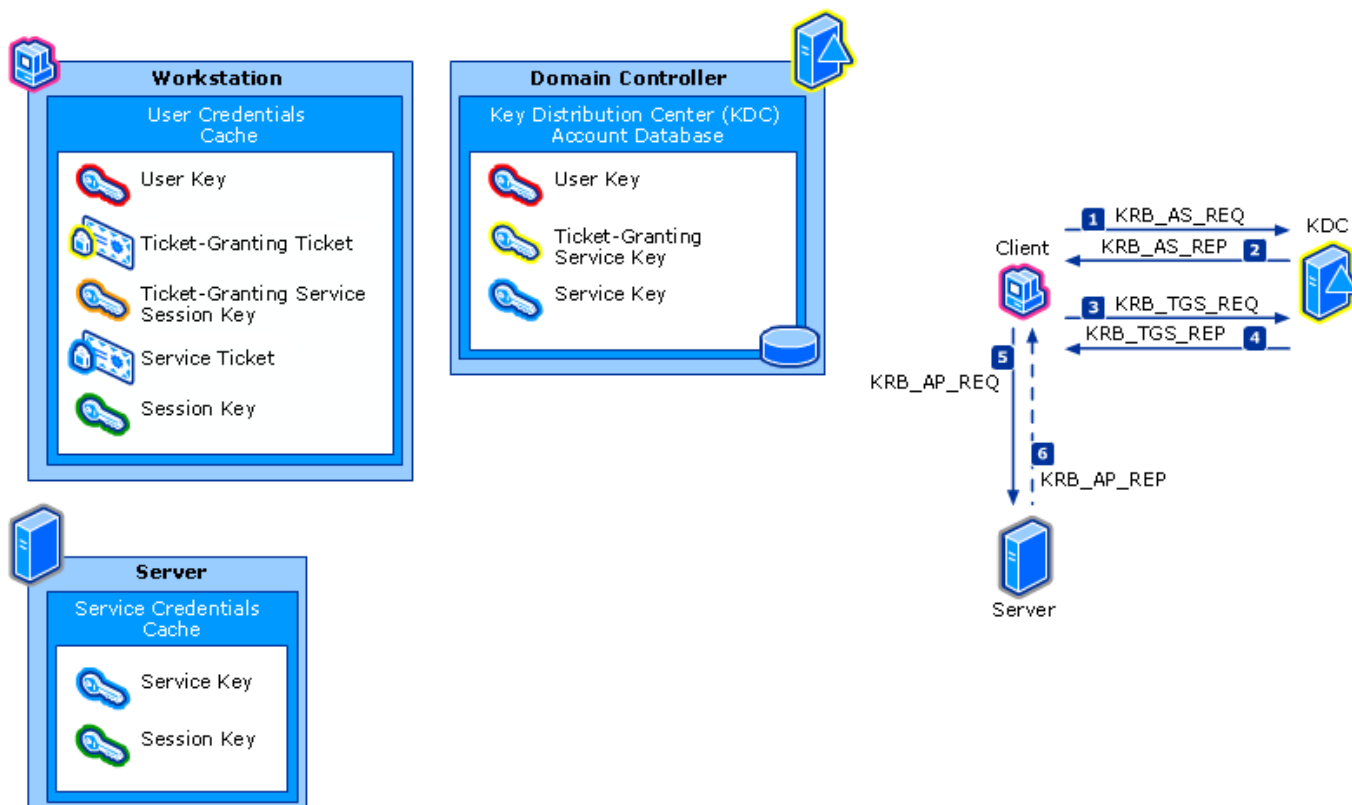
In one sentence: the client authenticates itself to the "authentication server" (AS), then demonstrates to the "ticket granting server" (TGS) that it's authorized to receive a ticket for a service (and receives it), then demonstrates to the "service server" (SS) that it has been approved to receive the service.



In more detail:

1. A user enters a username and password on the client.
2. The client performs a one-way hash on the entered password, and this becomes the secret key of the client.
3. The client sends a clear-text message to the AS requesting services on behalf of the user. Sample Message: "User XYZ would like to request services". Note: Neither the secret key nor the password is sent to the AS.
4. The AS checks to see if the client is in its database. If it is, the AS sends back the following two messages to the client:
 - Message A: Client/TGS session key encrypted using the secret key of the user.
 - Message B: Ticket-Granting Ticket (which includes the client ID, client network address, ticket validity period, and the client/TGS session key) encrypted using the secret key of the TGS.
5. Once the client receives messages A and B, it decrypts message A to obtain the client/TGS session key. This session key is used for further communications with TGS. (Note: The client cannot decrypt the Message B, as it is encrypted using TGS's secret key.) At this point, the client has enough information to authenticate itself to the TGS.
6. When requesting services, the client sends the following two messages to the TGS:
 - Message C: Composed of the Ticket-Granting Ticket from message B and the ID of the requested service.
 - Message D: Authenticator (which is composed of the client ID and the timestamp), encrypted using the client/TGS session key.
7. Upon receiving messages C and D, the TGS decrypts message D (Authenticator) using the client/TGS session key and sends the following two messages to the client:
 - Message E: Client-to-server ticket (which includes the client ID, client network address, validity period and Client/server session key) encrypted using the service's secret key.
 - Message F: Client/server session key encrypted with the client/TGS session key.
8. Upon receiving messages E and F from TGS, the client has enough information to authenticate itself to the SS. The client connects to the SS and sends the following two messages:
 - Message G: Client-to-server ticket, encrypted using service's secret key).
 - Message H: a new Authenticator, which includes the client ID, timestamp and is encrypted using client/server session key.
9. The SS decrypts the ticket using its own secret key and sends the following message to the client to confirm its true identity and willingness to serve the client:
 - Message I: the timestamp found in client's recent Authenticator plus 1, encrypted using the client/server session key.
10. The client decrypts the confirmation using its shared key with the server and checks whether the timestamp is correctly updated. If so, then the client can trust the server and can start issuing service requests to the server.
11. The server provides the requested services to the client.

3.1.2. MS kerberos usage



© 2003 Microsoft

- The Authentication Service Exchange

1. Kerberos authentication service request (KRB_AS_REQ)

The client contacts the Key Distribution Center's authentication service for a short-lived ticket (a message containing the client's identity and - for Windows client's - SIDs) called a ticket-granting ticket (TGT). This happens at logon.

2. Kerberos authentication service response (KRB_AS_REP)

The authentication service (AS) constructs the TGT and creates a session key the client can use to encrypt communication with the ticket-granting service (TGS). The TGT has a limited lifetime. At the point that the client has received the TGT, the client has not been granted access to any resources, even to resources on the local computer.

Why use a TGT? Couldn't the AS simply issue a ticket for the target server? Yes, but if the AS issued tickets directly, the user would have to enter a password for every new server/service connection. Issuing a TGT with a short lifespan (typically 10 hours) gives the user a valid ticket for the ticket-granting service, which in turn issues target-server tickets. The TGT's main benefit is that the user only has to enter a password once, at logon.

- The Ticket-Granting Service Exchange

3. Kerberos ticket-granting service request (KRB_TGS_REQ)

The client wants access to local and network resources. To gain access, the client sends a request to the TGS for a ticket for the local computer or some network server or service. This ticket is referred to as the service ticket or service ticket. To get the ticket, the client presents the TGT, an authenticator, and the name of the target server (the Server Principal Name or SPN).

4. Kerberos ticket-granting service response (KRB_TGS_REP)

The TGS examines the TGT and the authenticator. If these are acceptable, the TGS creates a service ticket. The client's identity is taken from the TGT and copied to the service ticket. Then the ticket is sent to the client.

The TGS cannot determine if the user will be able to get access to the target server. It simply returns a valid ticket. Authentication does not imply authorization.

- The Client/Server Exchange

5. Kerberos application server request (KRB_AP_REQ)

After the client has the service ticket, the client sends the ticket and a new authenticator to the target server, requesting access. The server will decrypt the ticket, validate the authenticator, and for Windows services, create an access token for the user based on the SIDs in the ticket.

6. Kerberos application server response (optional) (KRB_AP_REP)

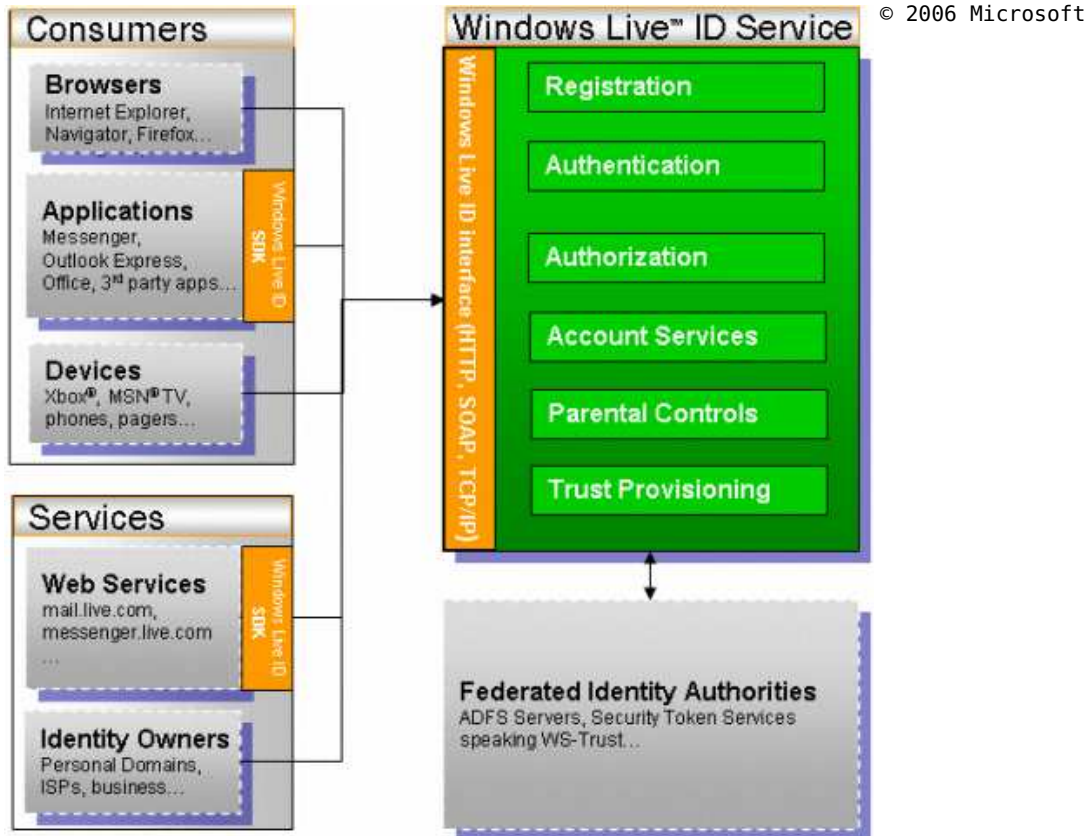
Optionally, the client might request that the target server verify its own identity. This is called mutual authentication. If mutual authentication is requested, the target server will take the client computer's timestamp from the authenticator, encrypt it with the session key the TGS provided for client-target server messages, and send it to the client.

3.2. Other SSO implementation examples

- Windows Live ID (ex .NET Passport)
- Liberty Alliance
- Google

3.2.1. Windows Live ID

Windows Live ID (originally named .NET Passport; briefly Microsoft Passport Network) is a "unified-login" service developed and provided by Microsoft that allows users to log in to many websites using one account. It was originally positioned as a single sign-on service for all web commerce.



3.2.2. WS-Trust

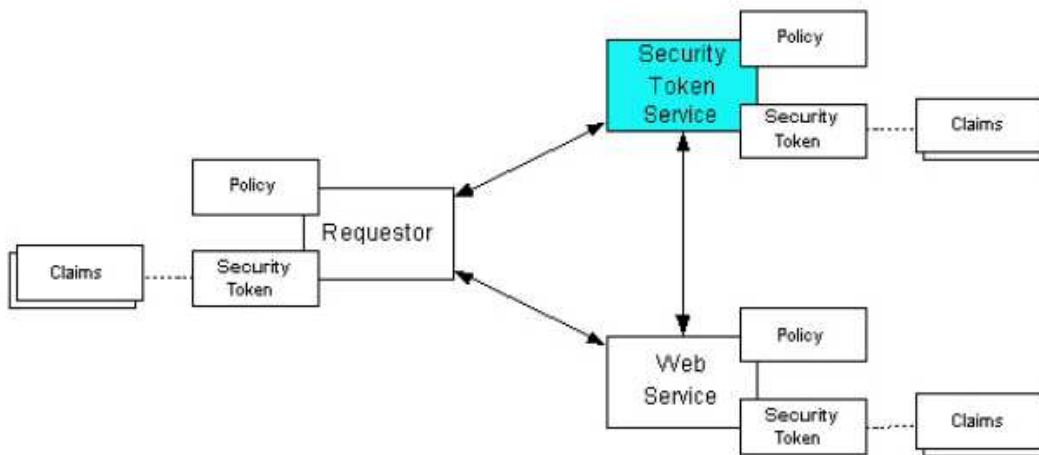
The Web service security model defined in WS-Trust is based on a process in which a Web service can require that an incoming message prove a set of claims (e.g., name, key, permission, capability, etc.). If a message arrives without having the required proof of claims, the service SHOULD ignore or reject the message. A service can indicate its required claims and related information in its policy as described by [WS-Policy] and [WS-PolicyAttachment] specifications.

Authentication of requests is based on a combination of optional network and transport-provided security and information (claims) proven in the message. Requesters can authenticate recipients using network and transport-provided security, claims proven in messages, and encryption of the request using a key known to the recipient.

One way to demonstrate authorized use of a security token is to include a digital signature using the associated secret key (from a proof-of-possession token). This allows a requester to prove a required set of claims by associating security tokens (e.g., PKIX, X.509 certificates) with the messages.

- If the requester does not have the necessary token(s) to prove required claims to a service, it can contact appropriate authorities (as indicated in the service's policy) and request the needed tokens with the proper claims. These "authorities", which we refer to as security token services, may in turn require their own set of claims for authenticating and authorizing the request for security tokens. Security token services form the basis of trust by issuing a range of security tokens that can be used to broker trust relationships between different trust domains.
- This specification also defines a general mechanism for multi-message exchanges during token acquisition. One example use of this is a challenge-response protocol that is also defined in this specification. This is used by a Web service for additional challenges to a requester to ensure message freshness and verification of authorized use of a security token.

This model is illustrated in the figure below, showing that any requester may also be a service, and that the Security Token Service is a Web service (that is, it may express policy and require security tokens).

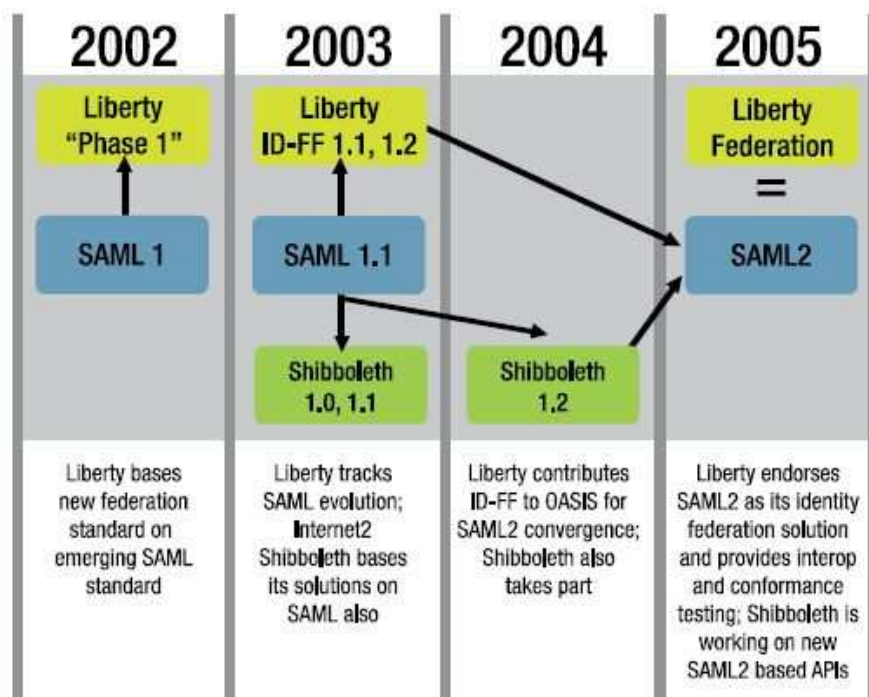


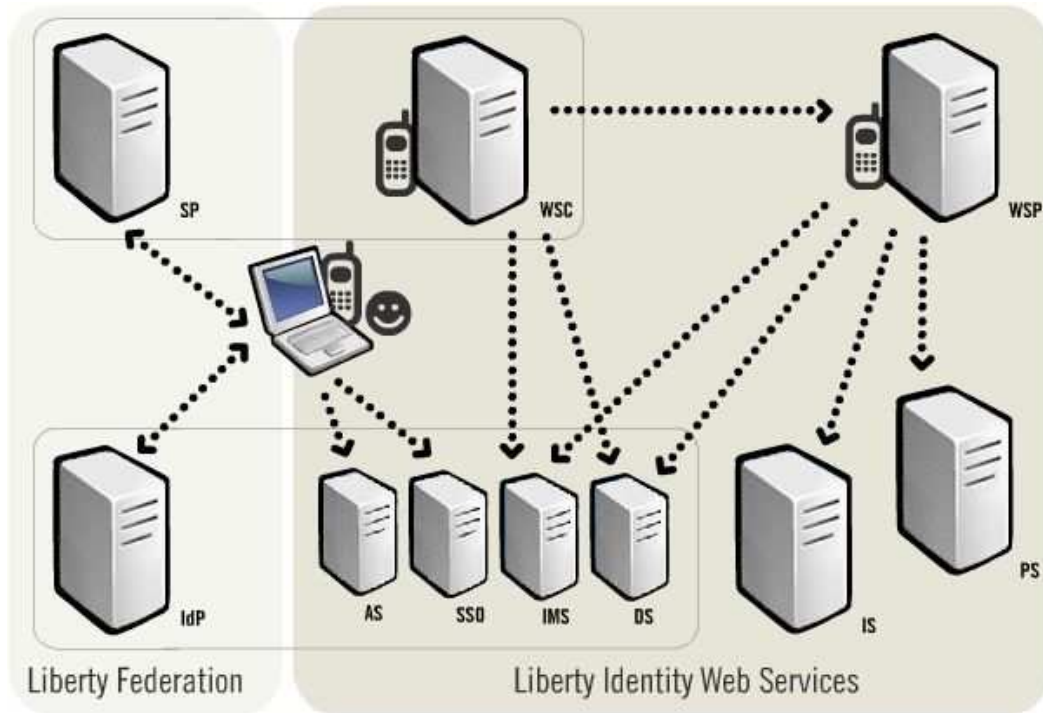
3.2.3. Liberty Alliance

The vision of Liberty Alliance is to enable a networked world based on open standards where consumers, citizens, businesses and governments can more easily conduct on-line transactions while protecting the privacy and security of identity information. This world, where devices and identities of all kinds are linked by federation and protected by universal strong authentication, is being built today with Liberty's open identity standards, business and deployment guidelines and best practices for managing privacy.

Members work closely together to:

- Build open standard-based specifications for federated identity and identity-based Web services.
- Drive global identity theft solutions.
- Provide interoperability testing.
- Offer a formal certification program for products utilizing Liberty specifications.
- Establish best practices, rules, liabilities, and business guidelines.
- Collaborate with other standards bodies, privacy advocates, and government groups.
- Address end user privacy and confidentiality issues.



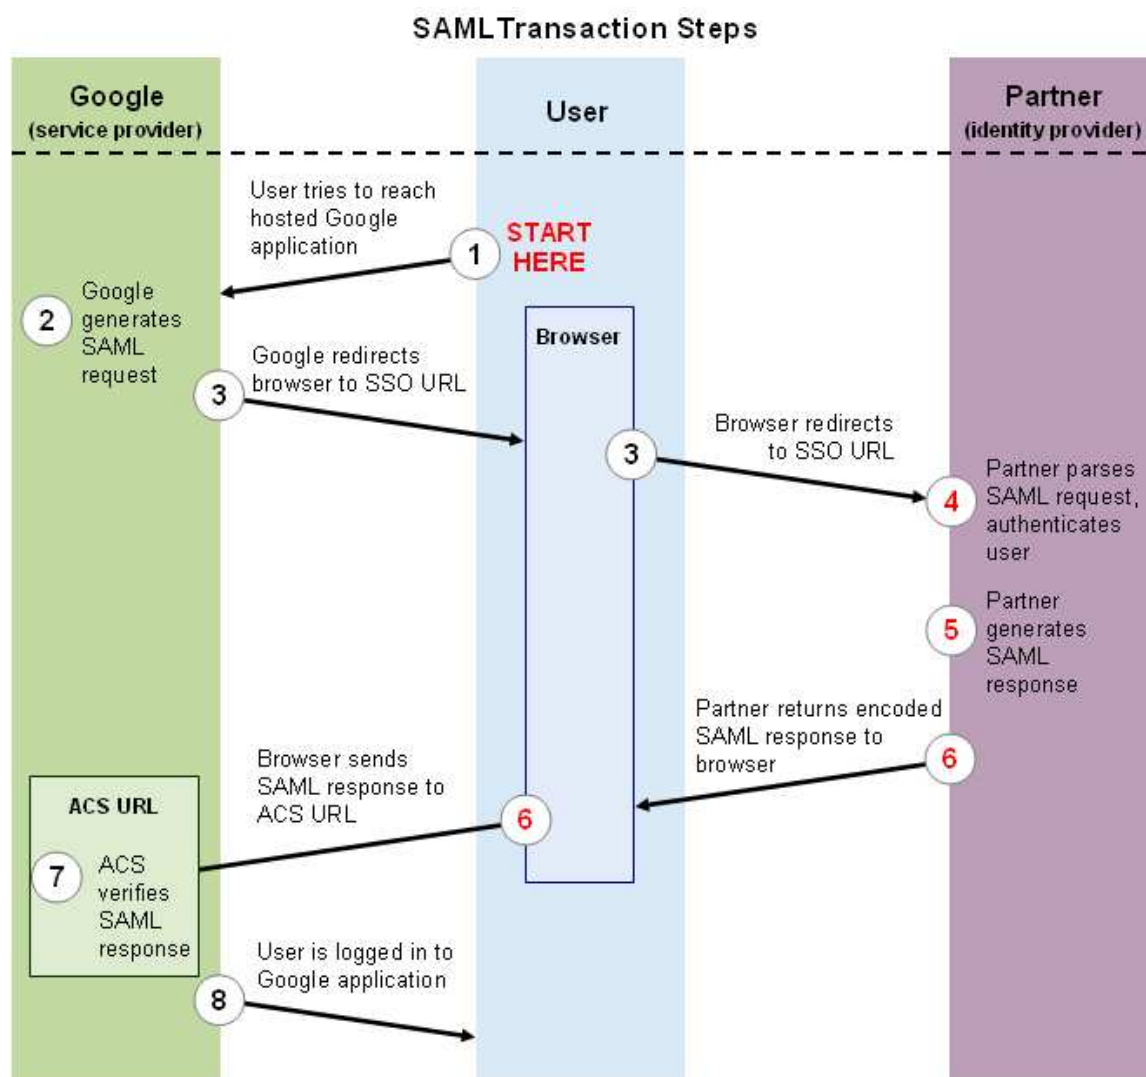


3.2.4. Google

Google Apps offers a SAML-based Single Sign-On (SSO) service that provides partner companies with full control over the authorization and authentication of hosted user accounts that can access web-based applications like Gmail or Google Calendar. Using the SAML model, Google acts as the service provider and provides services such as Gmail and Partner Start Pages (PSP). Google partners act as identity providers and control usernames, passwords and other information used to identify, authenticate and authorize users for web applications that Google hosts.

It is important to note that the SSO solution only applies to web applications. If you want to enable your users to access Google services with desktop clients such as Outlook - for example, Outlook would provide POP access to Gmail - you will still need to provide your users with usable passwords and synchronize those passwords with your internal user database using the Provisioning API.

The following process explains how a user logs into a hosted Google application through a partner-operated, SAML-based SSO service.



This image illustrates the following steps.

1. The user attempts to reach a hosted Google application, such as Gmail, Partner Start Pages (PSP) or another Google service.
2. Google generates a SAML authentication request. The SAML request is encoded and embedded into the URL for the partner's SSO service. The URL that the user is trying to reach is also embedded in the SSO URL.
3. Google sends a redirect to the user's browser. The redirect URL includes the encoded SAML authentication request that should be submitted to the partner's SSO service.
4. The partner decodes the SAML request and extracts the URL for both Google's ACS (Assertion Consumer Service) and the user's destination URL. The partner then authenticates the user. Partners could authenticate users by either asking for valid login credentials or by checking for valid session cookies.
5. The partner generates a SAML response that contains the authenticated user's username. In accordance with the SAML 2.0 specification, this response is digitally signed with the partner's public and private DSA/RSA keys.
6. The partner encodes the SAML response and the user's destination URL and returns that information to the user's browser. The partner provides a mechanism so that the browser can forward that information to Google's ACS. For example, the partner could embed the SAML response and destination URL in a form and provide a button that the user can click to submit the form to Google. The partner could also include JavaScript on the page that automatically submits the form to Google.
7. Google's ACS verifies the SAML response using the partner's public key. If the response is successfully verified, ACS redirects the user to the destination URL.
8. The user has been redirected to the destination URL and is logged in to Google Apps.

4. Identity management

The Laws of Identity

The "Laws of Identity" are intended to codify a set of fundamental principles to which any universally adopted, sustainable identity architecture must conform. The Laws were proposed, debated, and refined through a long-running, open, and continuing dialogue on the Internet. Taken together, the Laws define the architecture of the identity meta-system.

They are:

1. User Control and Consent

Identity systems must only reveal information identifying a user with the user's consent.

2. Minimal Disclosure for a Constrained Use

The identity system must disclose the least identifying information possible, as this is the most stable, long-term solution.

3. Justifiable Parties

Identity systems must be designed so the disclosure of identifying information is limited to parties having a necessary and justifiable place in a given identity relationship.

4. Directed Identity

A universal identity system must support both "omni-directional" identifiers for use by public entities and "uni-directional" identifiers for use by private entities, thus facilitating discovery while preventing unnecessary release of correlation handles.

5. Pluralism of Operators and Technologies

A universal identity solution must utilize and enable the interoperation of multiple identity technologies run by multiple identity providers.

6. Human Integration

Identity systems must define the human user to be a component of the distributed system, integrated through unambiguous human-machine communication mechanisms offering protection against identity attacks.

7. Consistent Experience Across Contexts

The unifying identity meta-system must guarantee its users a simple, consistent experience while enabling separation of contexts through multiple operators and technologies.

© Kim Cameron

5. Bibliographic references

- [RFC4301 Security Architecture for the Internet Protocol](#)
- [RFC4302 IP Authentication Header \(AH\)](#)
- [RFC4303 IP Encapsulating Security Payload \(ESP\)](#)
- [RFC4346 The Transport Layer Security \(TLS\) Protocol Version 1.1](#)
- [RFC3851 Secure/Multipurpose Internet Mail Extensions \(S/MIME\) Version 3.1 Message Specification](#)
- [RFC3852 Cryptographic Message Syntax \(CMS\)](#)
- [RFC3850 Secure/Multipurpose Internet Mail Extensions \(S/MIME\) Version 3.1 Certificate Handling](#)
- [RFC2631 Diffie-Hellman Key Agreement Method](#)
- [IETF/W3C XML-DSig Working Group](#)
- [Authentication \(Wikipedia\)](#)
- [The Failure of Two-Factor Authentication](#)
- [Kerberos \(Wikipedia\)](#)
- [Needham-Schroeder \(Wikipedia\)](#)
- [Kerberos page at MIT](#)
- [Designing an Authentication System: a Dialogue in Four Scenes](#)
- [How the Kerberos Version 5 Authentication Protocol Works in Windows Servers](#)
- [Single sign on \(Wikipedia\)](#)
- [SSO standardisation \(The OpenGroup\)](#)
- [Single Sign On \(AuthenticationWorld.com\)](#)
- [Windows Live ID \(Wikipedia\)](#)
- [Windows Live ID Service](#)
- [WS-Trust](#)
- [WS-Security \(Wikipedia\)](#)
- [Liberty Alliance \(Wikipedia\)](#)
- [Liberty Alliance](#)
- [SAML \(Wikipedia\)](#)
- [SAML Single Sign-On \(SSO\) Service for Google Apps](#)
- [Single Sign-On for TYPO3 \(and others\)](#)
- [OSSO \(Java Open Single Sign-On\)](#)
- [Laws of Identity](#)
- [identityblog](#)