Enabling Adaptive and Secure Extranets

Yves Roudier[1], Olivier Fouache[1], Pierre Vannel[2], and Refik Molva[1]
[1] *{roudier,fouache,molva@eurecom.edu}*
Institut Eurécom, 2229 route des Crêtes, B.P. 193, 06904 Sophia-Antipolis, France
[2] *{pierre.vannel@gemplus.com}*
Gemplus Labs Labs, Parc d'activités de Gémenos, B.P.100, 13881 Gémenos Cedex, France

Abstract: Extranets are tools that enable an organization to share part of its information system and infrastructure with other parties. Reaching this goal requires shielding from intruders while at the same time dynamically opening intranet resources. This article discusses how should such an extranet be designed. A solution that automates access control definition and enforcement is presented, which also addresses wide scale user management using a capability-based model. A prototype using the SPKI infrastructure is described that offers strong authentication thanks to smart cards.

Key words: extranets, firewalls, authentication, SPKI, Handle System, smart cards

1. INTRODUCTION

Today's business-to-business electronic commerce environments exhibit a major trait: information has to be shared with other parties such as suppliers, customers, partners, etc. On the other hand, totally opening the intranet is not acceptable: information is one of the main assets of a company, and thus needs to be thoroughly protected. Extranets are tools that enable a safe sharing. They interconnect the intranet of an organization with that of another party so that the organization can make its documents, services, computers, etc. available to a partner.

However, even though access control is becoming critical in a corporate extranet, the scale at which it must be managed does not make it easy for, say, two large organizations to interconnect. The number of available resources increases tremendously, as does the potential number of end-user

accessing them: this calls for the use of a fine-grained access control to resources. In addition to concerns about the granularity of access control, business relationships have to be implemented more frequently, sometimes almost immediately, and often for short periods of time: information rollout time has become a key factor for businesses. Finally, an intrusion in the corporate network is possible, be it from a hacker on the Internet or from an end-user of the partner organization.

The major problem with these requirements is that administrators cannot be expected to handle extranet security configuration at the frequency that is needed. In order to enable wide scale information sharing with other partners without compromising security, the security devices available in an extranet have to be automatically configured and operated, and in the most transparent way for end-users.

Is it possible to evolve the existing extranet paradigm to make it more suitable for B2B applications? This article proposes the *adaptive extranet paradigm* as a solution to these needs. In such a system, it is the automation and association of several traditional security tools like firewalls, public key certificate infrastructure, smart cards, and intrusion detection systems that enables a strong authentication of users and of their exchanges on a wide scale. This article also describes a prototype that was implemented along these guidelines in the SEVA project [SEVA].

Section 2 motivates the necessity of a new extranet paradigm based on the shortcomings of available technologies for setting up an intranet. Section 3 describes what should an adaptive extranet be made of. Section 4 details the definition and dissemination of application access rights. Finally, Section 5 explains the architecture chosen for enforcing these rights and gives some details about our prototype implementation.

2. TECHNOLOGIES FOR EXTRANETS

Corporate extranets have to provide access to a company resources. Although peer-to-peer technologies offer an attractive way to implement such an access [S-Peer], corporate networks are not ready yet. Many resources are closely tied with the server and service technology with which they were developed. This section considers several approaches to secure access to such resources.

2.1 Extranets vs. Virtual Private Networks

When it comes to securely interconnecting networks, the prevalent network architecture is that of virtual private networks. With this architecture, a

company's network can be extended beyond the boundaries of a firewall-controlled area. A virtual private network provides a network-level cryptographic traffic protection between two intranets. A good example of such setup is given by the IP Security Protocol [KeAt98a][KeAt98b], for instance in its tunnel mode. However, the VPN architecture implicitly defines the network as that of a single party only.

On the contrary, extranets are networks in which several partners share their resources and collaborate, as opposed to the VPN single party approach. In contrast with the VPN, which simply provides a tunnel for traffic, extranets must be defined from scratch in terms of the services offered. In particular, this implies defining an agreement on how these services can be used.

2.2 Authentication and Access Control with Firewalls

A first problem of the firewall architecture is that current security architectures are static: a human administrator is constantly required in order to update the filtering performed. Moreover, security entirely depends on the availability of this operator who must shut down any access of a malicious user. Finally this process is error-prone for filtering rules are complex to write.

Dynamic or adaptive firewalls are an attempt at solving this problem. For instance, the SunScreen firewall [SUN] introduces the so-called time of day rules, which are rules activated at a programmed hour. Other firewalls can be connected with an intrusion detection system. These firewalls can close a connection in case of an intrusion, but they do not solve the need for a dynamic configuration of extranets.

Current corporate network architectures generally associate firewall and server authentications: firewalls are first used to filter the traffic based on network-level elements, then users are authenticated on application servers with access control lists. LDAP databases help centralize these ACLs and share information with the firewall. In such an architecture, firewalls do not establish the corporate security perimeter anymore, but simply act as static malicious traffic filters. Application-level information is unfortunately unavailable to the firewall in order to adapt its filtering.

2.3 SOCKS

SOCKS [LGL96] is an IETF approved standard networking proxy protocol that enables hosts to access servers without requiring direct IP reachability. Authentication is possible and SOCKS might appear as a good solution to authentication on a firewall or more generally at the border of an intranet.

However, SOCKS requires modifying every communication call that is made in an client application. This is simply not possible with applications for which the source code is not available, that is the vast majority of applications in a corporate network. A new technology [eBorder] has been introduced recently that solves this problem. However, SOCKS authentication features are rather rudimentary compared to the needs of an extranet.

3. TOWARDS AN ADAPTIVE EXTRANET ARCHITECTURE

Local area networks, then intranets, still made it possible to identify a person in an organization and grant him one, sometimes several privileges for managing a file system or accessing an application server. The number of users envisioned in extranet architectures radically changed this point of view.

The use of capabilities, materialized by certificates is central here: no other solution makes it possible to manage inter-domain exchanges, that is, authenticating users from several companies dispersed over the Internet. The firewall configuration can be driven by the resolution of these capabilities.

We will call an extranet incorporating such mechanisms an adaptive extranet. The overall architecture of an adaptive extranet can be divided into three tiers, each of which has to handle a part of the complexity of securing business-to-business operations: the service management tier, the user management tier, and the extranet administration tier for configuring all security components.

3.1 Service Management Tier

An extranet consists of services offered by an organization to its partners. Establishing an extranet means deciding, at a high level, what services are available to a partner and being able to update this information. With the sheer volume of information on networks is exponentially growing, it appears that access control to the computers where this information is stored is not sufficient. The granularity of the access control should be smaller, but the multiform nature of digital documents (e.g. web pages, multimedia documents, applications…) as well as their organization make it impossible in practice to have a universal access control scheme.

Digital document references may provide an answer. The Corporation for National Research Initiatives (CNRI) has proposed the Handle System as “a general-purpose global name service enabling secure name resolution over

the Internet" [SRL01]. This proposal is currently under work at the IETF. It introduces a unique global namespace for digital resources over the Internet, with a root naming authority and sub-naming authorities, referencing all sub-naming authorities worldwide. An identifier to a digital resource is called a handle and can be resolved into a resource, for example an HTTP URL, or an FTP address.

Handles permit the resource administrator not only to define a uniform reference for accessing a resource, but more importantly, to introduce a very fine granularity of access control over the resources he provides. Access rights can be granted on each handle, independently from the application server and without modifying it. We view the handle system as an extensible framework for storing access control information, which makes it very suitable for defining multiparty application access rights.

The actual architecture of the Handle System, with a unique and public namespace, is not satisfying in the context of an adaptive extranet. The Handle System was thus modified in SEVA in order to introduce a private namespace per extranet. One extranet participant runs the extranet central authority to reference all naming authorities of the other partners (new feature). Every extranet participant runs a Local Handle Server (LHS). A handle on an extranet resource is only visible by extranet participants, an end-user being authenticated thanks to her smart card (new feature).

3.2 User Management Tier

Basically, the ultimate role of an authentication and access control architecture is to define if a user has been granted some right. This can be handled without too many hassles within a company's intranet. However, the several thousands of users from another company are managed by a different administrator than the one in charge of the resources available in the intranet accessed. Furthermore, managing an extranet is more complex: employees join and leave partner companies, partners join the extranet, etc. The successful deployment of any multiparty resource sharing architecture would imply that an important number of users would access a resource, even for a short moment: a capability-based access control model must be used in order to manage users' rights.

In addition, managing the users of another company directly would have privacy implications that a business would preferably avoid. Roles [SaSa97] represent a solution and can be introduced into the capabilities referred to above. In that respect, group certificates make the Simple Public Key Infrastructure (SPKI) ideal for managing the access rights granted over available resources of the extranet.

SPKI [EFL98a][EFL98b][EFL99] is destined to answer access control problems in wide-scale networks and is standardized by IETF: access control capabilities are encoded with authorization certificates. As opposed to X.509 [ITU88], SPKI is not designed for naming a party, but for stating its access rights. SPKI also permits the transitive delegation of access rights, which is essential for spanning the different jurisdictions of an extranet.

Handles offer an adapted resource designation for defining access rights and can be used straightforwardly in SPKI certificates. Handles can also make certificates smaller, for instance, they can reference the different access modes to a resource, which needs not be included in the certificate itself. Finally, handles make it simpler to lookup all the certificates required for proving that the requesting client has access to a resource.

3.3 Extranet Administration Tier

This tier lies in the middle of the two previous tiers and enforces the access control operations based on the information provided by those two tiers. In particular, it has to manage the firewall, and the user logs.

The administration tier also verifies that users have the right to access a resource. Adaptive extranets address applications where people from different corporate networks are accessing resources from another network and working together across the Internet. It thus becomes mandatory to design an architecture that not only protects from Internet hackers' attacks but also from malicious users even though their traffic is supposedly originating from a business partner. Users' rights are thus checked at the corporate system's border, as transparently as possible, then the behavior of authenticated users can be logged by the intrusion detection system. Suspect behaviors can entail the suppression of the user access to the local intranet.

The services available to every partner are also defined in an inital agreement that establishes the extranet. This agreement is used to generate the access rights of the allowed parties, and to configure firewalls and administration stations. XML [W3C00], or more precisely tpaML [tpaML] was selected to express such a contract: it can be used to automate the processing of the agreement that is updated each time a member joins or leaves the extranet. The rules regulating the extranet, like liabilities, security parameters, etc. can also evolve as well as the services offered and thus require establishing a new agreement.

4. ACCESS CONTROL DEFINITION AND MANAGEMENT

Access control definition is tremendously important in the adaptive extranet architecture. Access rights have to be distributed and stored first but they must also be managed. Access right management can be determined from the start thanks to an authorization certificate lifetime for instance, or might be needed suddenly in case of a network intrusion.

4.1 End-User Accreditation

The objective of the accreditation is to transcribe the extranet agreement within a participant's organization: that means deploying the extranet services for the authorized end-users and distributing them the corresponding access rights. Current organizations are hierarchical. The allocation of new services and their associated access rights follows a hierarchical path.

For instance, administrator stations rely on agents for automatically generating new SPKI certificates based on those that are regularly issued for every partner administrator. SPKI Delegation is heavily used in this process. At any hierarchical level, empowered end-users can delegate or restrict the access rights of their subordinates.

SPKI Certificates are very handy for avoiding revocation problems since they become automatically useless after their validity, which can be decided by the administrator, is over. The best way to exploit this feature is to manage certificate issuance through an automated infrastructure: such a system can be envisioned as a set of intelligent agents issuing a new certificate for every user according to the administrator settings. This part of our prototype is still under development based on the JADE agent platform [BPR00].

4.2 The Smart Card: a Management Tool for Access Rights

During the accreditation process, the end-user's smart card plays a key role. This smart card is a portable and personal wallet to access to the extranet services. It can be used to securely store cryptographic keys as well as the reference to services or resources. The smart cards used are Java Cards and offer multi-application support. They consist of a cryptographic card applet and of an embedded secure LDAP-like server [Mac00] where the service or resource references can be stored as well as access rights. The cryptographic card applet is a part of the authentication and access control scheme described in the next section.

The memory of current smart cards is still limited: only a few tens of kilobytes to store programs and data. The smart card is adequate to store secret cryptographic keys and to protect the cardholder's privacy. We thus chose to store only the reference to the certificate: the real certificate is stored into an LDAP server inside the intranet of the user's company rather than directly on the smart card.

4.3 Intrusion Detection and Access Right Management

The authentication of a user does not preclude the possibility of an attack on his part. The company offering the resource might in that case want to revoke the rights of a user performing repeated illicit or suspect operations.

Mobile users introduce yet another threat: these users can access services from a partner of their company or using a laptop computer from an Internet Service Provider, but they reside in a hostile environment. An attacker can more easily hijack their machine with a Trojan horse and then create havoc in the network of one extranet partner.

In both cases, a network intrusion detection system becomes an important part of the access control management infrastructure and these two elements should be interconnected. The snort [Roe] IDS has been integrated within our prototype adaptive extranet using the standardized IDMEF [CuDe01] exchange language. If an intrusion is detected, the liability of a user can be established with the help of the company who handles his identity and rights.

5. AUTHENTICATION IN THE EXTRANET

Once the access rights have been established by distributing certificates, actual traffic can be exchanged between the two intranets. Access control has to be enforced on this traffic to make it compliant with the access rights previously determined. An adaptive extranet must perform a strong authentication of the traffic exchanged, and in relation with the resources accessed.

5.1 An Authenticating Firewall

Access control must not only be described, but also enforced in conjunction with authentication. Firewalls [ChBe94] provide the most widespread technique for doing so and securing intranets: they enforce access control on the traffic by filtering packets and connections. This filtering security model is attractive because it is performed at intermediate network elements and thus leaves application servers unmodified. This

model is also quite easy to deploy, since all filtering operations can be performed in a centralized manner, for all servers of a domain for instance.

However, the most commonly used filtering model is that of the access control list, in which the operations authorized are listed for each user name. The architecture of an extranet requires filtering access based on capabilities instead as explained in Section 3.

5.2 Certificate Based Authentication and Access Control

Preserving performance is an important design guideline. This is why access control is enforced in two steps: application rights of a user are proven by sending a set of SPKI certificates; all traffic from that user is subsequently authenticated using lighter cryptographic operations.

As illustrated in Figure 1, whenever a handle (1) has to be resolved into an application resource (a URL in the example), the initial authentication is performed together with access control resolution (message (3)).

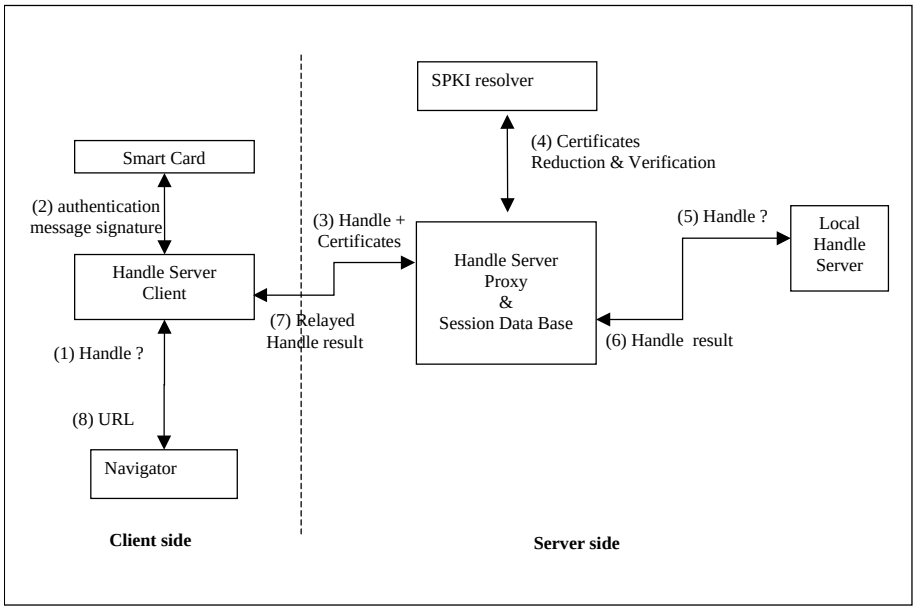

Figure 1. Verification of Access Rights

In our prototype, the Handle Server proxy was added to the set of proxies available in the TIS firewall toolkit [TIS]. It performs the associated SPKI certificate resolution (4) for proving the access rights of the end-user. It then also verifies that the message providing the handle is correctly signed by the smart card, which authenticates the user.

If the user is authenticated, the Handle Server proxy also performs the handle resolution (5) with a handle server modified to run without contacting the handle system root server hosted at CNRI. After this authentication, the user is logged in a session database and can subsequently be identified thanks to a session key transmitted in the authentication message (see section 5.3).

5.3 Traffic Tagging

The purpose of traffic tagging is to ensure that the packets transmitted were sent by the user, the sole entity owning the secret key, without having to pay the price of signing each and every packet transmitted. This has two important consequences: the traffic is strongly authenticated; a user can be logged and identified via his public key, although his name remains unknown to the firewall logging his accesses.

Tagging can be performed through the introduction of a specific communication layer (see Figure 2) in the client workstation. The data transmitted are encapsulated and marked with a cryptographic ticket (3). The ticket establishes the identity of the user based on the session key and ensures at the same time the integrity of the data: it simply consists of the keyed hashing of the data, using a session key established beforehand with the smart card. Encryption has not been provided but might as well be performed between intranets using the session key. However, it is not required for authentication only and we tried to limit the performance impact of the extranet security mechanisms.

A modified socket library intercepting any communication directed towards another SEVA intranet was programmed on Windows and experimented first (it offers a functionality similar to [eBorder]). The TCP packet was thus encapsulated with a specific traffic format, comprising the data, the authenticating cryptographic ticket, and some additional information like the destination address and port. With this network-level approach, it is possible to finely tune the granularity of authentication and thus the buffering of packets.

We also experimented with another technique: we focused on HTTP traffic, and finally implemented tagging via a simple HTTP proxy for its ease of deployment on several platforms (it currently runs on Linux or Windows). Using a proxy has the advantage of making it possible to use application-level information. For instance, instead of accessing resources through a handle, the user just types a URL in his browser and our prototype communication layer is able to convert this URL into one of the handles accessible to the user. Integration with applications becomes totally seamless.

The ticket verification is performed on the firewall by the corresponding proxy (4), in relation with a previously authenticated handle resolution. This proxy verifies that the ticket was constructed correctly. The proxy also decapsulates the original traffic before relaying the connection (5).

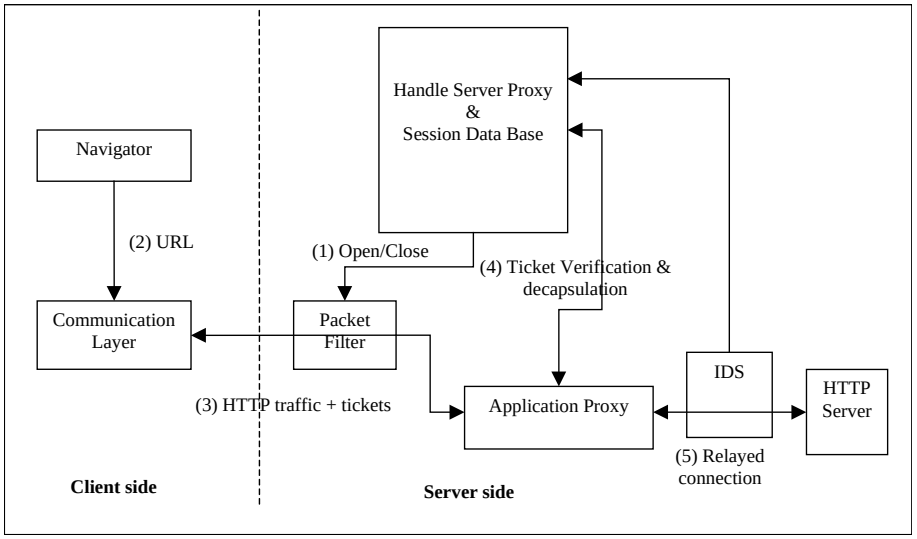

Figure 2. Traffic Tagging

The smart card is used here as a real-world key to the adaptive extranet: it keeps the private key of the user that does not reside on the workstation itself and is used as an essential element in a zero-knowledge authentication protocol. No traffic with another partner can take place on the extranet without the smart card because the ticketing process cannot run short of a valid session key. Only the smart card can generate a session key for authenticating the traffic. Inserting a card in a workstation reader enables the workstation to connect during a fixed time slot. At the end of this slot, a new session key is automatically reestablished with the handle service proxy. For performance reasons however, ticketing cannot performed on the card, which can only exchange protocol data units with a serial line.

6. CONCLUSION

Business-to-business corporate services require the secure interconnection of the networks of different parties. This interconnection is

difficult because on one hand, services and the resources that they reference are complex, and on the other hand, identity and end-user's rights can only be defined and handled within the user's company. Furthermore, B2B relations necessitate a very swift configuration of security devices whose frequency cannot be handled by human operators. The security devices used in extranets presently do not address these requirements. Adding security devices like an intrusion detection system to an extranet will not solve these problems *per se* either.

This article proposed to have a certain number of otherwise classical security devices collaborate in a coordinated fashion to extend the capabilities of a corporate firewall in order to realize what can be called an adaptive extranet architecture. Users and access rights are central to such an adaptive extranet and must be handled separately: both can be managed using a SPKI public key infrastructure and a generic notion of resource named a handle. The user's traffic is authenticated between his workstation and the firewall of the corporate network accessed with a lightweight cryptographic scheme so as not to degrade access performance. The tight integration of SPKI authentication with the traffic ticketing process makes it possible to achieve the automatic configuration of the firewall. Smart cards play an important role for authentication as well and are used instead of a password. They are personal security devices that integrate perfectly with the automatic configuration of the extranet security devices thanks to the SPKI model: they enable users to prove his identity or administrators to issue authorization certificates. Finally, even though strong authentication of end-users is enforced, the system does not threaten end-user privacy.

An adaptive extranet prototype has been developed and deployed by the SEVA project team that already implements all security features for access control and certificate automated distribution. Results thus far are very encouraging: the access to a web server can be totally controlled without servers being modified at all by the deployment of this system. The SEVA prototype implementation also demonstrates that client applications can be retrofitted with strong authentication without any modification. We experimented with the introduction of access control for UDP traffic as well, and would be interested to support it more completely in the future.

ACKNOWLEDGEMENT

SEVA is a collaborative project of ATOS, Electricité de France, Eurécom, and Gemplus, and is supported by the French Ministry of Economy, Finances, and Industry and the Réseau National de Recherche en Télécommunications (RNRT).

REFERENCES

[BPR00] F. Bellifemine, A. Poggi, G. Rimassa, and P. Turci. *An Object Oriented Framework to Realize Agent Systems*. in Proc. of WOA 2000 Workshop, Parma, May 2000, pp. 52-57.

[ChBe94] B. Cheswick and S. Bellovin. *Firewalls and Internet Security: Repelling the Wily Hacker*. Addison Wesley, 1994, ISBN 0-201-63357-4

[CuDe01] David Curry, Hervé Debar. *Intrusion Detection Message Exchange Format Data Model and Extensible Markup Language (XML) Document Type Definition* - draft-ietf-idwg-idmef-xml-06.txt. December 2001

[eBorder] http://www.permeotechnologies.com/technology/wpapers.htm. Permeo Technologies. *e-Border white papers*.

[EFL98a] C. M. Ellison, B. Frantz, B. Lampson, R. Rivest, B. M. Thomas, and T. Ylönen. *Simple Public Key Certificate*, Internet draft <draft-ietf-spki-cert-structure-05.txt>, March 1998.

[EFL98b] C. M. Ellison, B. Frantz, B. Lampson, R. Rivest, B. M. Thomas, and T. Ylönen. *SPKI Examples*, Internet draft <draft-ietf-spki-cert-examples-01.txt>, March 1998.

[EFL99] C. M. Ellison, B. Frantz, B. Lampson, R. Rivest, B. M. Thomas, and T. Ylönen. *SPKI Certificate Theory*, RFC 2693, September 1999.

[ITU88] ITU-T. *Recommendation X.509: The Directory - Authentication Framework*, 1988.

[KeAt98a] S. Kent, R. Atkinson. *IP Authentication Header (RFC 2402).* November 1998.

[KeAt98b] S. Kent, R. Atkinson. *IP Encapsulating Security Payload (ESP) (RFC 2406).* November 1998.

[LGL96] M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, L. Jones. *RFC 1928. SOCKS Protocol Version 5.* March 1996

[Mac00] A. Macaire. *An Open Terminal Infrastructure for Personal Services.* TOOLS Europe 2000, 5-8 June 2000, Le Mont-St-Michel, France

[Roe] Martin Roesch. *Snort - Lightweight Intrusion Detection for Networks* - http://www.snort.org/docs/lisapaper.txt

[SaSa97] R. S. Sandhu, P. Samarati. *Authentication, Access Controls, and Intrusion Detection*, in The Computer Science and Engineering Handbook, pp 1929-1948, 1997.

[SEVA] SEVA project home page - http://www.eurecom.fr/~nsteam/SEVA/

[S-Peer] Texar. *S-Peer*. http://www.s-peer.com/

[SRL] S.X. Sun, S. Reilly, L. Lannom. *Handle System Namespace and Service Definition.* IETF Draft. May 2001.

[SUN] SUN Microsystems. SunScreen Secure Net 3.1, Technical Whitepaper

[TIS] FWTK.ORG *unofficial page on TIS firewall toolkit* - http://www.fwtk.org/main.html

[tpaML] IBM. *Electronic Trading-parter Agreement for E-Commerce. ebXML proposed specification, version 1.06.* http://www.ebxml.org/project_teams/trade_partner/

[W3C00] World Wide Web Consortium. *Extensible Markup Language (XML) 1.0. W3C Recommandation.* http://www.w3.org/TR/2000/REC-xml-20001006