

Authorization & Endpoint Security

The Bachelor Series

Jericho in depth...

Authorization & Endpoint Security

Securing clients in Jericho networks

Leon Teheux

Capgemini's Security & Innovation Research Centre, based in the Netherlands, focuses on near future IT Security solutions. The Jericho forum's vision on network de-perimeterization and Boundaryless Information flow™ has been the starting point for this research centre. Research papers from this centre appear in two distinct categories;

1. The Master Series

Researcher holding a masters degree in Informatics or are in the process of obtaining a master degree publish in the Master Series. The participating University and the Capgemini Security & Innovation Research centre have approved publications in this category.

Publications in this Series for 2008;

- Jericho in depth... Secure Communications by A. Stan
- Jericho in depth... The road to Jericho by A. Stan

Planned publications in this Series for 2008;

- Demystifying trust by F. van Leijden
- Jericho in depth... Automated Security Classification by K. Clark
- Jericho in depth... Trust Management for Trust brokers by A. Demarteau

2. The Bachelor Series

Researchers holding a bachelors degree in Informatics or in the process of obtaining a bachelors degree publish in the bachelor series. Their University and the Capgemini Security & Innovation Research centre have approved publications in this category.

Publication in this series for 2008;

- Jericho in depth... Endpoint security by L. Teheux
- Jericho in depth... Authentication and Accounting by E. Barannikov
- Jericho in depth... Trust broker Services by A. Bruning
- Jericho in depth... Trust broker framework by A. Bruning

Planned publications in this series for 2008;

- Jericho in depth... Controlling the COA framework by J. Willemsen
- Jericho in depth... Fully ASP based by D. Hanenberg & F. Aardoom

Copyright © Capgemini 2008

All rights reserved. No part of this work may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without the prior written permission of Capgemini.

Preface

Since the dawn of the Internet at the ending of 1969 a lot has changed, I'm sure nobody will disagree with a statement like that. During the last couple of years however, we seem to have hit a mid-life crisis of the Internet. The sudden boost of Internet technology over the past decade does not fit well with our outdated design principles for network security. Most organizations hold tight to their fortress approach in trying to protect the internal network from the hostile Internet. Understandable, but not really realistic. In the Netherlands, we are particularly proud of our water management techniques. In a country that lays for more then sixty percent below sea level we know that we have to build and maintain solid dikes to prevent our country from flooding. Having holes in these dikes quickly diminishes the whole purpose of have a dike. The same holds true for perimeter defence in computer networks. Information leakage via email, hyves, my space or mobile data solutions like iPod or USB diminishes the purpose of perimeter security. Today's business world is one of collaboration, one of working together., one of global markets. The Internet is the ideal candidate to support this collaboration. The Jericho Forum (Open Group), formed by Security professionals from the largest organisations in the world described their vision of network de-perimeterization and boundryless Information flow™ in various publications. These visions formed the starting point for Capgemini's Security & Innovation Research Centre.

Together with the best universities in the Netherlands, Capgemini's offers academic researchers and graduate students to ability to conduct empirical academic research into the topic of Collaboration Oriented Architectures or to conduct feasibility studies into the Jericho Forums visions.

Marco Plas

Head of Jericho Research
Capgemini Security & Innovation Research Centre
Capgemini Netherlands

Contents

Preface	5
Introduction	8
1 Aims of the study	9
2 The Jericho Forum	10
3 The Jericho Forum Model	16
4 Endpoint Security in the Jericho Forum Model	21
5 Authorization in the Jericho Forum Model	48
6 Conclusion and Future Research	60
7 References	62
Appendix Protocols and Standards	63

Introduction

During the past few years, network security has been steadily advancing as a “hot item” on the lists of IT managers. Both software- and hardware vendors have been developing new and improved products to enhance the security of existing networks. The result of this has been the creation of virtual bastions, where accessing and sharing data has become increasingly difficult. When comparing modern business models to current network architecture, a multitude of incompatibilities can be observed. Contrary to the transparent way of conducting business, a modern network can best be compared to a castle surrounded by stone walls and moats. Information has to transit multiple boundaries in order to reach its destination. This architecture does not allow businesses to fully leverage network possibilities or seamlessly exchange information with its partners. The Open Group’s Jericho Forum has proposed a unique and elegant approach to this problem: instead of trying to secure a network, the network itself should be secure. This process of de-perimeterization can be achieved by using inherently secure protocols, configurations and systems. At this time no formal specifications of this process have been defined. Capgemini, being one of the Open Group’s Platinum members, has assembled a team of open minded and enthusiastic young professionals to research possibilities and develop solutions. Capgemini’s Jericho Forum Research Group has divided this challenge into several interacting and interdependent processes. This document describes the aspects of Authorization and Endpoint Security. This document will firstly introduce the Jericho Forum and establish an overview of project directions. These directions are translated into processes, after which this document will expand upon the Authorization and Endpoint Security processes. Based upon information gained from contact with vendors and available literature, it will explain the requirements, establish process interactions, compare current implementations with the requirements and, finally, make recommendations on how to proceed.

Leon Teheux

1. Aims of the study

The intention of the Jericho Forum Research Project is to define requirements and solutions that can be used to create an implementable network based upon existing and proven technologies. The Research Project has been split into several directives. This paper will research and describe the Endpoint Security and Authorization processes. In order to achieve these intentions, this study has been written with the following aims:

- Endpoint Security:
 - To determine the role of Endpoint Security within the Jericho Forum model
 - To determine logical requirements
 - To determine interaction requirements with other processes
 - To determine technical requirements
 - To compare the established requirements to currently available solutions
 - To recommend currently available solutions
 - To recommend steps to be taken for future implementations
- Authorization:
 - To determine the role of Authorization within the Jericho Forum model;
 - To determine logical requirements;
 - To determine a logical solution;
 - To determine interaction requirements with other processes;
 - To recommend steps to be taken.

2. The Jericho Forum

The Jericho Forum Research Group

The Jericho Forum has published several whitepapers and position papers regarding various subjects it considers important to an open network. However, at this time¹, no document with a suggested architecture has been published.

In order to determine the shape of the research project, a general overview of a network based upon the Jericho Forum vision had to be created. Based upon the Vision Whitepapers and the 11 Commandments the following general requirements were established:

- Security mechanisms must be pervasive and scalable;
- All devices must be capable of maintaining their security policy on untrusted networks;
- All people, processes and technologies must have been authenticated and transparent levels of trust are necessary for any transaction to take place;
- Mutual trust assurance levels must be determinable;
- Authentication, authorization and accounting must interoperate with other implementations outside your area of control;
- Access to data should be controlled by security attributes on the data itself.

When combining these requirements with the subjects established by the published whitepapers, the following project research directions can be proposed:

- Federated identity;
- Trust relationships;
- Open and inherently secure protocols;
- Endpoint security;
- Digital rights management.

In order to provide insight into these research directions and how these intermix, a short description of each of these is provided below.


Federated identity

As established by the Federated Identity vision paper, the majority of today's user authentication schemes still depend on the use of a username and password to identify users. The burden on users of managing large numbers of username and passwords has led to proposals for Federated Identity systems, where a single set of credentials can be used to authenticate with several organizations, which have agreed to work together as a federation. The Federated Identity approach has been proposed as a business-to-business service for employees, where one

¹ June 15, 2007

organization manages the user credentials and authorization to systems run by another organization.

An example of such a scheme is depicted in the image below. An independent third party provider is responsible for the account information used by multiple companies.



At the moment, most approaches have been limited to authenticating human users. However, within the Jericho model, devices, applications and resources also need to be able to authenticate themselves.

Challenges

The establishment of a common scheme of data attributes that can be requested is essential. In addition, in order to comply with the Jericho Forum Commandments, methods that allow a trusted environment to be established by mutual and peer to peer authentication using open and secure protocols should be researched.

Relation with Endpoint Security

As described above, endpoint devices also need to be able to authenticate themselves. This means that they will need to be able to provide proof of identity, which, combined with a status established by the Endpoint Security process, will determine their effective permissions.

Relation with Authorization

The intention of Federated Identity is to let third parties establish and maintain identities and authorizations. As such, the implementation of Federated Identity may determine the area of responsibility of the Authorization process.


Trust relationships

Business transactions require a level of trust between participants, which, in this context, means that each partner has confidence that the other will fulfill his part of a bargain. Trust, in the business context, relies primarily upon contracts to specify the behavior that is required and an enforcement mechanism to punish and deter non-performance. For this to function in an online environment, a process is required that registers and verifies the identity of each party involved. However, registration processes are hard to automate and, therefore, expensive. This prevents businesses from implementing trust based relationships with other companies.

Challenges

The cost of registration can largely be reduced by sharing it between the parties involved. This is facilitated by mechanisms such as federation, which is designed to share identities and authorizations between organizations, thus extending their use. Current federation mechanisms are orientated towards federating customer identity between members of a supply chain and, if agreed upon, between related supply chains. They aim to facilitate interactions between a customer and an organization. However, organizations require additional functionalities. The federation process needs to be easier to automate and should allow the creation of new mechanisms and functionalities, such as reputation, for sharing trust information. In addition, a common legal infrastructure is required to facilitate and support a de-perimeterized online environment.

A linking pin between companies and information systems that can fulfill these challenges is needed. A Trust broker service, or Trust broker, that is trusted by all parties involved can meet this need.



Relation with Endpoint Security

As stated in the eighth Jericho Forum Commandment, authentication information obtained within one company must be able to be exchanged with other companies. As such, Endpoint Security results must be able to be used outside the scope of a company. Therefore, Trust broker services should be able to include these results and exchange them.

Relation with Authorization

The establishment of trust relationships between entities has a direct impact on the authorization process. Depending on contracts and agreements signed, authorization decisions may be made.

Open and inherently secure protocols

With nearly every enterprise using computers that regularly connect to the Internet, businesses employing wireless communications internally and the majority of users connecting to services outside the enterprise perimeter, de-perimeterization of networks has become a reality. The fourth Jericho Forum Commandment states that the use of inherently secure and open protocols is essential to provide protection from insecure data transport environments. Ideally, secure protocols should act as fundamental building blocks for secure distributed systems, being adaptable to the needs of applications whilst adhering to requirements for security, trust and performance.

Challenges

In order for inherently secure protocols to become adopted as standards, they must be open and interoperable. The Jericho Forum has stated that before a protocol can be universally adopted, it should be fully open, royalty free and well documented.

Relation with Endpoint Security

When communicating with verification servers and other entities, endpoint devices need to communicate their status in a secure manner, using protocols determined by this project direction.

Relation with Authorization

Authorization protocols should use secure protocols to prevent certain attacks, such as man-in-the-middle or identity theft, from happening. In addition, open protocols are required to ensure interoperability between systems developed by different vendors.


Endpoint Security

A number of recent trends, such as home and mobile working, de-perimeterization, outsourcing and the growth of e-commerce, have resulted in a growth of the number of devices involved in transactions. These are placing new demands on business' capability to provide trusted access to their services. As described by the Jericho Forum Commandments security solutions should be

simple, pervasive, scalable and be able to provide mutual trust assurance. The intention of Endpoint Security is to provide the ability to raise the level of inherent trust in computing devices to a point where all devices involved in a transaction meet the criteria of trust for that transaction.

Current endpoint security solutions are generally limited to validating clients trying to connect to an environment, with the trust being one-way; the client not being able to form an opinion regarding the environment it is connecting to. One-way trust may lead to attacks such as phishing, where attackers attempt to gather credentials by impersonating trusted entities. Being able to mutually establish the trust level of end points allows more valuable transactions to take place electronically.

Having devices or users from multiple organizations validate their trust level upon trying to interact with your information, as opposed to validation when connecting to your network zone, enables more flexible and secure ways of working.



Traditional NAC

Challenges

Standards are required in order for security agents placed on endpoints to be able to interoperate, and for end points to require only a single agent. This allows agents to expand onto a wide variety of end points such as phones, PDAs, network devices and PCs. Companies should not completely depend upon endpoint agents for security, network behavior analysis should be implemented to discover unusual traffic and prevent attacks from taking place.

Relation with Authorization

Depending on the amount of authorization needed by accessed information, the Authorization process may need to request additional parameters, which may be established by the Endpoint Security process, in order to determine applicable rights.

Digital Rights Management

Within the Jericho forum mode, digital rights management is especially focused on the classification of data and the management of rights to access that data. The DRM process is not limited to the classification of text documents; it is intended to be able to classify data ranging from telephone calls to video to databases and it should be able to operate on any device that is able to access information.

Challenges

Jericho Forum Commandment number nine states that access to data should be controlled by security attributes on the data itself, whilst Commandment number eleven states that data must be appropriately secured when stored, in transit and in use. At this moment, even on systems under your control, the storage of insecure and typically unencrypted data that is reliant on system or even network security controls is flawed. A lost PC with client information or a database administrator who has access to all personal information in a database are both examples of where the data is inappropriately protected. Especially on data held outside of your immediate control, there is a need to be able to manage, change or revoke access, as well as the need to manage versioning and to reduce concurrency.

Relation with Endpoint Security

Digital rights management is responsible for maintaining data access rights. However, information leakage due to compromised systems may be outside the scope of DRM. The Endpoint Security process can prevent these leaks by detecting compromised devices and prohibiting them from accessing stored data.

Relation with Authorization

The Digital Rights Management process is responsible for classifying data and determining rights applicable to data. These rights are used by the Authorization process to determine if an entity can access that data and what needs to be known about the accessing entity before interactions are allowed.

3. The Jericho Forum Model

In order to unify the aforementioned program directions into a single interoperable network, the directions need to be translated into processes. These processes are responsible for delivering predefined input and output, upon which actions and decisions can be taken. By combining the IETF² AAA-framework, as described in RFC 2903 and 2904, with the program directions established, a network matching the requirements set forth in the Jericho Forums whitepapers and commandments can be designed. In order to allow future modifications to the model or implementations, a modular design is required. To do so, all processes have strictly defined areas of operation and responsibility, with standard methods of interaction between them.

The following project directions were initially established:

- Federated Identity
- Trust relationships
- Open and inherently secure protocols
- Endpoint Security
- Digital Rights Management

The AAA-framework is used to implement the Federated Identity and Trust relationships project directions. Open and inherently secure protocols and their implementation are the responsibility of the Encryption process. Endpoint and general network security have become the responsibility of the Endpoint Security process. Finally, the Digital rights management project direction has been translated into the Data Classification process.

In conclusion, a Jericho Forum based network consists of the following processes:

- Authentication
- Endpoint Security
- Authorization
- Accounting
- Data Classification
- Encryption

In order to understand process responsibilities, a short description of these processes, comparing them to project directions is given below. Each process is further explained in their respective papers.

² Internet Engineering Task Force

Processes

Authentication

As described within the Federated identity vision paper, entities on the network should be able to mutually determine each others identities. Within the AAA-framework, this process is known as authentication. At the moment, many different methods of authentication exist. Entities on a network may identify themselves by supplying a username and password, by having a unique digital key known as a certificate, by providing biometric identification or any combination of these possibilities. All these methods have one thing in common: they require the existence of databases storing user credential information. In most cases, these databases are only accessible by users within a certain domain. For example, a user may want to log on to his company PC, access his Hotmail account and complete a transaction on E-bay by means of online banking. This requires the use of 4 separate identities. The eighth commandment of Jericho states: *“There must be the capability of trusting an organization, which can authenticate individuals or groups, thus eliminating the need to create separate identities.”* This commandment requires a different method of authentication. A possible authentication method that matches the commandments that require the mutual assurance of trust levels and an interarea interaction of the authentication and authorization process is known as the claim-based authentication process. Within the claim-based authentication process, entities are no longer recognized by the information they provide that matches an internal database. Instead, entities are expected to provide claims to the resource they are accessing. Consider the following example: when buying alcohol you are required to provide proof of age. You can provide a claim by showing the salesperson your passport or driving license. In this context, the official papers are your claims and the salesperson trusts your claims, since they come from a recognized and trusted source. This providing of proof is analogous to the claim-based authentication process. By providing enough claims from a trusted resource (in this case the government), identities can be established. Although a claims-based architecture may be very compatible to the Jericho Forum model, it should be determined whether it is possible to implement such an architecture with existing solutions, or whether alternative architectures should be considered.

The Authentication process has been described by Evgeny Barannikov in his book, *Jericho in Depth – Authentication and Accounting*, Capgemini, 2008.

Endpoint Security

As described within the Endpoint security direction, the Endpoint security process is responsible for providing the means to establish inherent trust levels between endpoints, with the intent to create a situation where all the devices involved in a transaction meet the criteria of trust for that transaction. At the moment, many Endpoint security, or Network Access Control, solutions exist. However, most of these solutions were not designed to interoperate with other solutions and they

lack the ability to verify all network devices³. Most solutions only provide Endpoint Security for personal computers running certain operating systems. Several of the Jericho Forum commandments refer to the Endpoint Security process. The second commandment of Jericho states that “Security mechanisms must be pervasive, simple, scalable & easy to manage”, whilst the fifth commandment states that “*All devices must be capable of maintaining their security policy on an untrusted network*”. The seventh commandment states: “*Mutual trust assurance levels must be determinable*”. These commandments require a solution where every device connected to a network should be able to participate in the Endpoint security process. This means that a universal standard should exist that governs agent behavior and interactions. In order for a secure device to function in a possibly insecure network, devices must be able to maintain their security policies. This means a solution should exist that can monitor device status and act upon it, essentially requiring agents installed on a device.

Authorization


As described within the Federated Identity vision paper, after establishing identities, a process should be in place that determines what rights are applicable to requests. Within the AAA-framework, this process is known as authorization. At the moment, authorization is dependent on the authentication process. Only after an entity has established its credentials, the authorization process can determine what rights are associated with the information it attempts to access and, acting upon this, allows or denies the request. Contrary to existing models, within the Jericho Forum model authorization is not limited to merely making decisions. It can actively require other processes to send additional information in order to make more informed decisions. This ability to communicate with other processes enables creation of a flexible and dynamic authorization process.

Accounting

The final part of the AAA-framework is accounting. Accounting refers to the tracking of actions and events. For each defined interaction a log entry is created. Traditionally, these log entries could be used for multiple purposes, for example billing based upon the number of transactions, evidence in court or as information for auditing. One of the main purposes of the traditional accounting model is the ability to provide information after facts have occurred. Within the Jericho Forum model, accounting should be used in a completely different way. Accounting stores information regarding user interactions. These interactions may include information regarding the location the user logs on from and the device the network is accessed from. Instead of passively storing this information, it can be used to determine a user's default behavior. This process is known as baselining and is already commonly used in network behavior analysis and credit card transaction processing. After establishing an initial baseline of user behavior, new user interactions can be compared against this baseline and a value can be attached to the interaction in progress. Based upon the expected behavior and the behavior in progress, actions may be taken. Automatic measures, such as

³ Gartner MarketScope for NAC, 2007

blocking access or informing management, may be taken when important information is accessed in a manner not included in the baseline. In order to make this baseline available for external sources, but maintain confidentiality, a network entity called a Trust Repository should be established. This entity is responsible for maintaining user baselines from multiple sources, whilst ensuring privacy and accuracy.



The Accounting process has been described by Evgeny Barannikov in his book, *Jericho in depth – Authentication and Accounting*, Capgemini, 2008.

Data Classification

As described within the Enterprise Information Protection & Control (Digital Rights Management) position paper, this process is responsible for the management of rights to access digital data, the control over that data, the usage of the data and the integrity of the data. In order to do so, it is necessary to be able to determine the importance of the data relative to the policies of the company and apply relevant options. Data in this context refers to all methods in which information may be captured, be it in text documents, sound files or movie files. The Data Classification process should analyze the content of documents and combine this with metadata attached to documents in order to determine the relative importance of the information to a company. Based upon this relative importance, certain options may be applied to the data, for example the necessity that certain authentication attributes are included within the authorization process before it is allowed to be accessed or the need to use certain encryption options. Not only should the Data Classification process assign certain rights, it must make sure that when data importance changes due to a lapse in time or unveiling of a product, access rights are changed to reflect this. In order to provide enterprise-spanning information classification, methods should be established that allow all information processing devices within an enterprise to determine data relevance and rights when processing information and verify or update

rights attached. In order for data to be able to cross organizational borders whilst maintaining the ability to determine access rights, the data itself should be able to determine what rights are applied to it. This aspect of Digital Rights Management is an essential part of the Data Classification process.

The Data Classification process will be described by Kas Clark in his book, *The Jericho Forum Project – Data Classification*, soon to be published.

Encryption

As described within the Protocols position paper and referenced to by other position papers, the Encryption process is responsible for establishing parameters and options to ensure secure communications and data storage. The intention of the Encryption process is to establish which combination of protocols and options may be used in order to secure information in all its transitional stages. This means different options may be available for data in storage, transfer and processing. In order to determine what security or encryption solution is required, several factors have to be taken into account. These factors are established by security policies, combined with information received from the Data Classification process and intercompany relationships.

The Encryption process has been described by Alina Stan in her book, *Jericho in depth – Secure Communication*, Capgemini, 2008.

4. Endpoint Security in the Jericho Forum Model

In the previous chapter, the processes that form a network based upon the Jericho Forum model were described. In addition, the relative position of Endpoint Security within this model has been established. The goal of this chapter is to establish Endpoint security process requirements and interactions. In order to do so, a clear and expansive examination of the available documents describing Endpoint Security as envisioned by the Open Group's Jericho Forum is required. There are two documents available that must be examined: the Endpoint Security vision paper and the Eleven Commandments. After examining these documents, it will be possible to establish the requirements needed, the way the Endpoint Security process will interact with other processes within the model and to identify available solutions and see how they compare to the established requirements.

The Eleven Commandments

The Eleven Commandments is the official document published by the Open Group's Jericho Forum, which defines both the areas and the principles that must be observed when planning for a de-perimeterized future. The commandments serve as a benchmark by which concepts, solutions, standards and systems can be assessed and measured. This chapter describes these commandments and describes the relation between them and the Endpoint Security process. Based upon these commandments, logical requirements will be established. This dissection is based upon version 1.2 of the vision paper, published May 2007.

1. The scope and level of protection should be specific and appropriate to the asset at risk.

The implemented Endpoint Security solution should be cost effective and it should not implement security measures in areas where they are not relevant.

2. Security mechanisms must be pervasive, simple, scalable and easy to manage.

The Endpoint Security process should be able to protect all devices on the network, whilst not excessively increasing user and administration interactions.

3. Assume context at your peril.

The Endpoint Security solution should be able to operate in different environments, although some configuration flexibility should be allowed.

4. Devices and applications must communicate using open, secure protocols.

In order for the Endpoint Security solution to become universally implemented, the protocols used should be open in order to provide for peer assessment and trust among implementers.

5. All devices must be capable of maintaining their security policy on an untrusted network.

The Endpoint Security solution should be able to connect to any network and maintain a secure state, as defined by its security policies.

6. All people, processes and technology must have been declared and transparent levels of trust must exist for any transaction to take place.

This commandment requires that all entities involved in a transaction are aware of their obligations and that trust requirements may vary on the transaction in question. The Endpoint Security solution should provide different metrics that can be used to establish endpoint integrity, based upon transactional requirements.

7. Mutual trust assurance levels must be determinable.

The Endpoint Security solution should be able to establish mutual trust in order to prevent attacks such as Phishing.

8. Authentication, authorization and accountability must interoperate with solutions outside your area of control.

The Endpoint Security solution process should be able to exchange information with entities outside the companies' infrastructure, either directly or through Trust broker services.

9. Access to data should be controlled by security attributes on the data itself.

Data security attributes may require the use of an operational firewall on the host data access is requested from. The Endpoint Security solution should be able to answers these demands.

10. Data privacy requires a segregation of duties/privileges.

It should not be possible for users to modify their Endpoint Security agent settings. As such, rights relating to the Endpoint Security solution should be reserved for a separate account, only used for configuring the solution.

11. By default, data must be appropriately secured when stored, in transit and in use.

Data must always be appropriately secured. This means that when the Endpoint Security process transmits or generates data, a suitable solution should be used to maintain integrity and confidentiality.

The Endpoint Security vision paper

The Endpoint Security vision paper is the official document published by the Open Group's Jericho Forum that describes the endpoint security problem as envisioned by the Jericho Forum, a recommended solution to this problem and what steps should be taken in the near future. This chapter contains the essentials of the vision paper and expands upon certain areas. Based upon these essentials, technical requirements will be established. This dissection is based upon version 1.0 of the vision paper, published October 2006.

Problem:

- There is a need to provide trusted access to services.
 - During the MediaPlaza meeting of April 18th, 2007, a discussion was sparked by the question whether or not banks are to be responsible for the security of customer devices. Although the representatives of both the Rabobank and ABN-AMRO bank acknowledged the benefits of this vision, they claimed current technologies were not accurate and encompassing enough.
- Endpoint Security should raise the level of inherent trust.
 - As stated in the Commandments, Endpoint Security should encompass all devices on a network and therefore be able to determine what hosts are compromised, bypassing them in favor of secure devices.
 - Depending on the sensitivity of the data, determined by the Data Classification process, trust level and options may vary.

Why should I care?

- Endpoint Security must establish mutual trust.
 - Mutual trust allows safer transactions, enabling more sensitive and valuable transactions to take place.
 - According to statistics from the Anti-Phishing Working Group⁴, phishing attacks are on the rise. Phishers attempt to acquire sensitive information, such as usernames, by masquerading as a trustworthy entity in an electronic communication. By establishing mutual trust, phishing attacks can be eliminated.

⁴ Anti-Phishing Working Group, January 2007 Phishing Activity Trends

Current concept

- Endpoint Security operates by managing endpoints and network security zones.
 - Endpoints can be all network devices (hubs, switches), access devices (workstations and mobile devices) and servers.
 - A security zone is group of devices networked under a common security contract together. Appropriate security policies must be applied for each zone.
 - Zones may include all types of devices. For example, all remote devices accessing a network may be part of a zone, it being irrelevant if it is a PDA or laptop.
- An endpoint's security posture is any security attribute for the end point that a remote party may wish to rely upon.
 - Examples of security postures are an up to date firewall or virus scanner.

Flaws in the current concept

- In general, traffic can only enter and leave a zone through a zone security device. Resulting in:
 - Establishing a single point of failure;
 - Making this point susceptible to DOS attacks.
- Special agents are required on devices not compliant with the IEEE 802.1X standard, increasing administration efforts and costs.
- Interorganizational usage of traditional clients is difficult because of a lack of standards.
 - Although interoperability is not a necessity in traditional networks, it is mandated by the Jericho Forum Commandments.
- On-demand installation of agents is unlikely to work in situations where the end point is locked-down.

Requirements for enhanced solutions

- An organization needs the capability to register end points from many sources, for example its own, customers' and suppliers' endpoints.
- Endpoints need a capability to be registered in several organizational zones simultaneously.
- For systems that interact using inherently secure protocols, both systems must be capable of validating the trust.

Challenges to the industry

- Standards are required so that single agents placed on endpoints can interoperate.
- Standards are required for bi-directionally secure communication.
- Collaboration is required to develop a secure protocol that will allow a security agent on an endpoint to validate remote endpoints.

Requirements

There are 2 different kinds of requirements that are used to determine the architecture desired. The logical requirements are intended to create a general overview of the Endpoint security process and focus on establishing process flows.

The technical requirements are intended to determine what characteristics a protocol or solution should have and as such are intended to support the logical requirements.

Logical Requirements

Based upon the 11 Commandments, four logical categories can be established: Trust, Security, Scope and Manageability. Trust describes the Endpoint Security requirements necessary for the establishment of a global trust infrastructure. In order for data to follow secure paths across a network, all relevant devices it passes through must be able to verify their security status.

All relevant devices in this context refer to any device that:

- a) is running any kind of operating system or software that can be compromised
- and
- b) has access to the unencrypted data.

The verification of their security status means that any device meeting the previous criteria must be able to deliver certain claims, depending on those required by the authorization process, to the Trust broker service. In addition, mutual trusts are necessary to prevent data being lost by identity theft. Security describes the requirements necessary for the establishment of a secure, global network. The primary objective is the establishment of secure endpoints, which means that any device meeting the criteria stated in the Trust section must be able to maintain their security policies and report on their condition, regardless of where and how they are connected to a network. The Scope describes the requirements necessary for the establishment of a global network. This requires endpoints to be able to communicate certain aspects of their status to entities outside the local area of control, whilst maintaining user privacy. The establishment of a global network requires a scalable infrastructure, both in size and scope. The final category is Manageability. This category describes the requirements necessary for the establishment of a manageable network. Manageability has two equally important requirements: Endpoint Security should be focused and should be managed as close to the device as possible. This means that Endpoint solutions should –if possible- be installed on the device containing data, and be managed by entities that are closely related to the data stored, for example the users themselves. This requires an easy to manage Endpoint solution.

Summarized, the following logical requirements are established:

Trust

- endpoint security must determine trust levels for all relevant devices
 - endpoint security must span all tiers of the architecture
 - endpoint security must be able to support mutual trust

Security

- Endpoint security must be able to protect the device it is operating on.

Scope

- Endpoint security must be able to operate with a global scope.
 - Endpoint security must be scalable.

Manageability

- Endpoint security should be simple to manage.
- Endpoint security should be managed as close to the device as possible.

Technical Requirements

Based upon the Logical Requirements and the Endpoint Security whitepaper, technical requirements can be established. These can be divided in the following categories: Operation, Protocols, Management and Standards. Operation refers to the manner in which implementations should act. First of all, Endpoint Security Agents must be able to operate on any relevant device. As described within the Logical Requirements, this refers to any device that is running any kind of operating system or software that can be compromised and has access to the unencrypted data. Furthermore, Agents must be able to communicate with any other Agent, regardless of the developer. Communications are not restricted to the exchange of endpoint verification status; it may include messages indicating that the remote entity is not allowed to request a verification. Finally, to prevent network outages, the Endpoint security process must not introduce any Single Point of Failures (SPOF) to the network. The Protocols category refers to the requirements that protocols must meet. As described within the Endpoint Security whitepaper, protocols used should be open and inherently secure. In addition, the implementation of protocols should be transparent in order to be able to replace certain elements of a solution without needing to replace the entire solution, thus future-proofing the network. The Management category refers to the management options the solution should offer. Important is the ability to define what external sources are allowed to request host verification. In order to prevent administrators from changing policies on certain systems, segregation of duties is an important aspect of security management. Finally, the Standards category demands that all protocols used must adhere to open standards. The use of open protocols promotes peer-review and enhances the interoperability between solutions.

Summarized, the following technical requirements are established:

Operation

- Agents must be able to operate on all relevant devices.
- Agents must be able to communicate with any other Agent.
 - Endpoint status may be delivered in claims.
- No single point of failure must exist.

Protocols

- Endpoint security should use secure protocols.
- Protocols used must be transparent and be able to be replaced.

Management

- Segregation of duties should be implemented.
- It must be possible to apply rights to external devices accessing endpoints under your control.

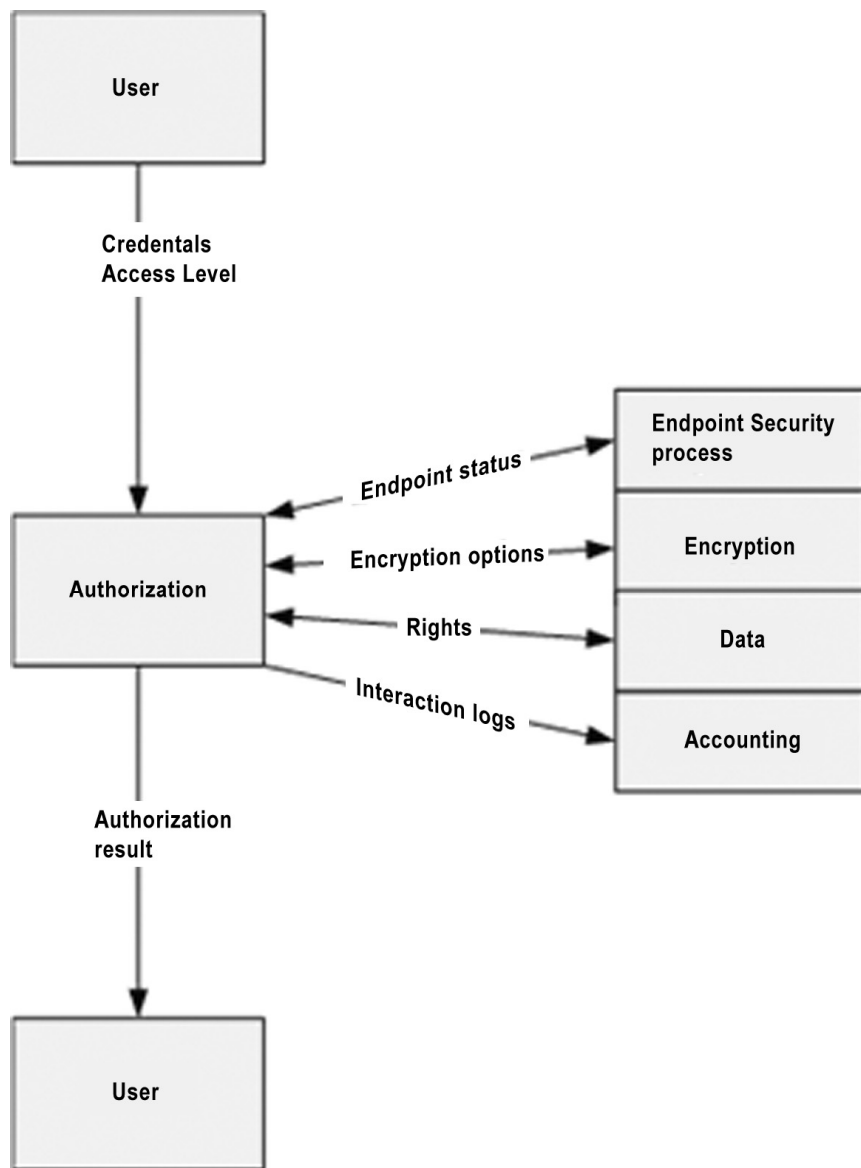
Standards

- All protocols and implementations should adhere to open standards.

Process Interactions

Endpoint Security interactions

In order to provide a modular framework, interactions with other processes in the Jericho Forum model should be formalized. This enables standardized communications between modules, improving compatibility with other solutions. The first step is to determine the overall process interactions. This requires the establishment of the process flow. As visualized in image 4.3.3.1, the Authorization process has to send a request to the Endpoint Security process for verification of the host status. After receiving this request, the Endpoint Security request interacts with the Encryption process to establish cryptographic protocols and options to be used. After verifying the device, the Endpoint Security process returns the results to the Authorization process. All actions performed in this flow must be able to be logged by the Accounting process. An interaction with this process is therefore required.



This can be translated to the following inputs and outputs:

Input

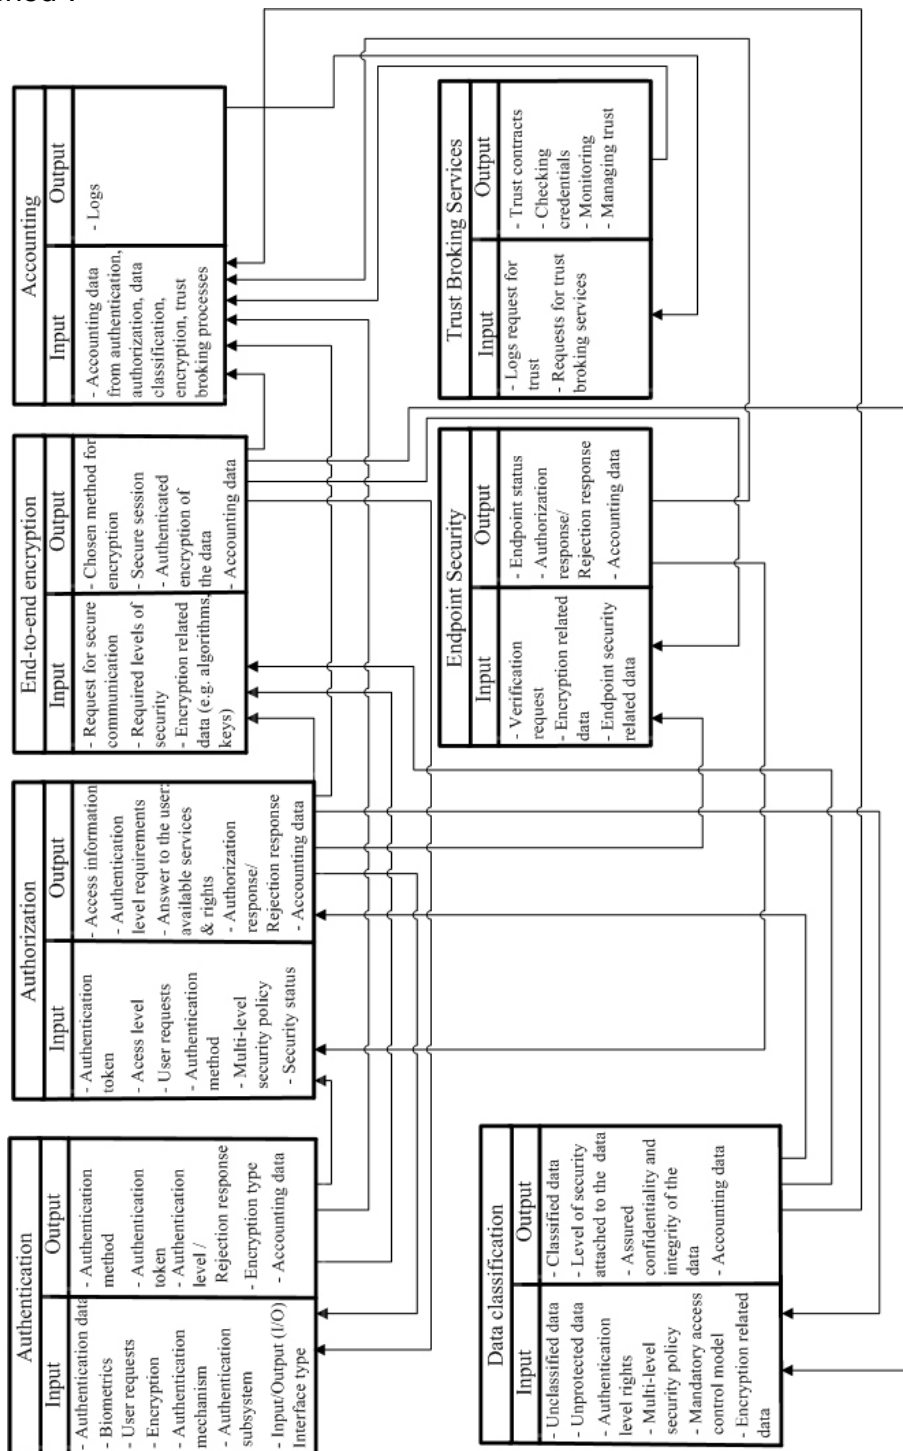
- Authorization (request for verification)
- Encryption (encryption options to be applied)

Output

- Authorization (verification status)
- Accounting (listing of actions)
- Encryption (request for encryption options)

Jericho Forum Network Interactions

When combining the interactions established by other processes with the Endpoint Security process, the following overview of process interactions can be established⁵:



Jericho Forum interactions

There are several interesting observations that can be made. Important to note is the connection between all processes and the Accounting process. This arrangement ensures that all network events can be monitored from a central location. This enables the creation of a Trust Repository where accounting information can be combined to create a baseline of interactions.

Logical Solutions

In order for Endpoint Security to become an integral part of the Jericho Forum model, it must be universally available and implemented. This means that an architecture will have to be developed that will allow any endpoint to request the status of any other endpoint, whilst maintaining the privacy of both devices. When considering the established Requirements and the Endpoint Security whitepaper, two approaches to a global Endpoint Security network exist. The first approach is the use of a secure, peer-to-peer protocol that will allow Endpoints to mutually exchange verification information. The second approach promotes the use of a trusted third party, a Trust broker, to gather Endpoint Security status information. Both approaches are discussed below.

Peer-to-peer Architecture

In traditional client-server networks, clients are able to connect to dedicated servers that contain data or provide processing power. In order to provide more flexible and survivable connections, the peer-to-peer architecture abolishes the idea of network clients and servers. Within a peer-to-peer architecture, network entities can directly connect to each other and serve each others requests.

There are several advantages to Peer-to-peer networks:

- all nodes on the network are part of the system, eliminating any single points of failure in the system, increasing availability
- all nodes provide their own bandwidth and resources, thus increasing accessibility
- as no additional hardware is required, a lower cost

Disadvantages of Peer-to-peer networks include:

- a decentralized architecture inherently means lack of administrative control of Endpoint rights and rules
- additional administrative effort may be required to incorporate systems within the Accounting process
- all endpoints must communicate using the same protocol
- a decrease in user privacy, as all endpoints can request information from any other endpoint


In order to negate some of these disadvantages, the Peer-to-peer architecture has been modified to include some level of centralization. These modifications implied that the Pure peer-to-peer model has evolved into another model: the Hybrid peer-to-peer model.

This means the following peer-to-peer models exist:

Pure peer-to-peer

- Peers act as equals, merging the roles of clients and server
- There is no central server managing the network
- There is no central router


This image shows how computers interact on a pure peer-to-peer network:



Hybrid peer-to-peer:

- Has a central server that keeps information on peers and responds to requests for that information
- Peers are responsible for hosting available information (as the central server does not have it), for letting the central server know what information they want to share, and for making its shareable information available to peers that request it

The image below shows how the first request of the host gets answered by a central server. This server determines whether or not the request is valid and if all request options are acceptable. After establishing the validity of the request and the information available to be shared, it returns a token to the initiator of the session. The initiator can then present the token to the desired endpoint and receive the information requested.



Trust broker Architecture

Being a client-server model, the Trust broker Architecture is dependent on multiple servers that communicate with each other in order to provide Endpoint Security status for endpoints. A key element within this architecture is the existence of Trust brokers that are responsible for determining and communicating the status of endpoints under their control.

There are several advantages to a Trust broker based architecture:


- The centralized architecture means administrative control remains within the organization endpoints belong to.
- The ability to customize applications, as the organization is completely in charge of its local Endpoint Security solution.
- A centralized architecture decreases the amount of network routing needed, increasing responsiveness of applications.

Disadvantages of a Trust broker based architecture include:

- A decrease in availability as the Trust broker responsible for intercompany communications may introduce a single point of failure;
- An increase of nodes or usage of the network may slow down the network, decreasing accessibility.

- A new infrastructure supporting the Trust broker is required.

The image below shows how host verification is handled within a Trust broker Architecture. First, an endpoint sends a request for verification to its local Trust broker. This Trust broker then queries the Trust broker in charge of the remote endpoint, which in turn verifies the Endpoint status and replies with the verification status.



Recommended Logical Solution

Traditionally, solutions are rated upon three factors that encompass all aspects of its operation: Availability, Scalability and Manageability. In order to be able to recommend a logical solution, all proposed architectures will be compared to these factors, within the framework established by the Jericho Forum.

Availability

Availability refers to the amount of time a service is available and the level of service offered. High availability means that a certain service is almost always usable, whilst the level of service refers to a predetermined performance baseline that a service should meet. When financial transactions are dependant on the results of a verification, a model providing for a very high Endpoint verification availability is necessary.

Within the Pure Peer-to-peer model, endpoints can query another endpoint directly. Since no other endpoints are involved, no single point of failure exists. Availability is only influenced by the status of the network connections and the system status of the endpoints attempting to communicate.

Within the Hybrid Peer-to-peer model, endpoints query a server which returns a response or token, containing authorization to the requested information. Clients then connect to the endpoint and request a verification by transmitting the token. The server containing an index of endpoint information may be a single point of failure. If it is possible for endpoints to switch to a Pure Peer-to-peer model if the main server becomes unreachable, availability may be maintained.

Within the Trust broker Architecture, endpoints send verification requests to a central server. That server then becomes responsible for the verification process. This server is essential for a functional Endpoint Security process and, unless proper precautions are taken, may represent a single point of failure.

Scalability

Scalability refers to the ability of the service to gracefully handle increasing amounts of load. The Jericho Forum states that the Endpoint Security solution must be able to operate on a global scale. An extremely scalable model must be implemented to provide such a solution. Since each endpoint must be able to deliver a verification report when requested, care should be taken that either the resources available can meet heavy loads or that short-term caching of results can be applied.

Within the Pure Peer-to-peer model, each endpoint can directly query the endpoint to be verified. Since each endpoint provides its own resources (e.g. bandwidth and CPU time), few limitations regarding resources exist. However, since each host must be able to answer queries, all endpoints on the network should run an agent that conforms to the standard at that time, increasing cost of management. No true technical limitation exist within a Pure Peer-to-peer model, however, operational costs may prove to be prohibitive.

Within the Hybrid Peer-to-peer model, endpoints query a server which then determines the request options to be allowed and returns a token to the initiator. The initiator can then deliver the token to the endpoint to be verified and receive the information requested. The introduction of a server in the Pure Peer-to-peer model may limit scalability by requiring increasing amounts of server resources. In addition, by making servers responsible for providing tokens, conflicts regarding area of responsibility may decrease scalability.

Within the Trust broker Architecture, endpoints become the responsibility of a predetermined server. When a verification is required, the endpoint transmits a request containing request options to the server, which then becomes responsible for providing a result. This server will determine which server is

responsible for the endpoint to be verified and will send a request to this server, which then authorizes and runs the scan and finally returns the results to the initial Trust broker.

The dependency on central servers for communications means that Endpoint Security services are limited to the resources available to servers. Since each server becomes responsible for providing Endpoint Security services for devices in its domain, an infrastructure based upon the model of the existing Domain Name System (DNS)-style architecture⁶, which faced similar requirements and challenges, may provide scalability. In addition, since communication is limited to conversations between Trust brokers, various organizations may use different and customized agents, with only Trust brokers needing to adhere to standards.

Manageability

Manageability refers to the ability of the service to be controlled. Within the Jericho Forum model, this includes the ability of administrators or users to control which entities are allowed to ask for verification. In addition, management tasks should be performed as close to the data as possible.

Within the Pure Peer-to-peer model, all Endpoint Security network interactions are authorized and performed on the endpoints themselves. This means that policies can only be applied on the endpoint that receives the request to be verified. Since no central entity is required, accounting is performed by the endpoints themselves. Management tasks are placed in the hands of the end-users. However, this means that organizational security policies may become harder to enforce.

Within the Hybrid Peer-to-peer model, policies can be applied on both the server that replies to the initial request and the endpoint itself. Accounting should be performed on the endpoints, with additional information being generated by the server. Management tasks are performed by users and IT staff.

Within the Trust broker Architecture, multiple policing points exist. Both local and remote policies can be applied to requests and the logging of actions can be centralized. In addition, customized agents can be placed on user workstations, increasing integration with other enterprise requirements. Within this model, management tasks can be performed by users and IT staff.

Conclusion

The preferred model should provide the optimum combination of Availability, Scalability and Manageability. The architectures are ranked based upon how well

6 Libor Dostalek & Alena Kabelova, *Dns in Action*, Packt Publishing Ltd, 2006

they conform to the requirements, with 1 being the most matching solution. The solution with the lowest total is to be preferred.

	Pure Peer-to-peer	Hybrid Peer-to-peer	Trust broker Architecture
Availability	1	2	3
Scalability	3	2	1
Manageability	3	2	1
<i>Total</i>	<i>7</i>	<i>6</i>	<i>5</i>

With a score of 5, the Trust broker Architecture is the recommended solution to the requirements set forth by the Jericho Forums commandments. The Trust broker architecture has the additional advantage of being able to integrate with other Jericho Forum processes that use a similar architecture.

Technical Solution

As discussed within the Aims of the Study, the intention of the Capgemini Jericho Forum Research Group is to deliver a model that can be implemented using existing solutions. In order to meet this goal, the characteristics of several existing solutions will be compared to the established logical and technical requirements. Explained below are several key factors that determine the suitability of existing products. Based upon how well an implementation matches the logical and technical solutions, a recommendation will be made.

Mode of operation

The mode of operation refers to how a solution integrates with the network. It may use pre-installed clients, dissolvable clients, no clients at all or a combination of these possibilities. The Jericho Forum requires a solution that can verify any device on a network, preferably by using a common agent.

Network integration

Network integration refers to the impact a solution has on network architecture. For example, it may require network devices to be reconfigured, appliances to be installed inline or only work with equipment of certain vendors. The Jericho Forum requires a solution that has a minimum impact on existing network infrastructures.

Resource requirements

Resource requirements refer to the impact a solution has on network resources. Certain implementations may be less suitable for slower connections or may require intensive processing on endpoints. The Jericho Forum requires a solution that can be used on any endpoint. Therefore, resource usage must be as low as possible.

Device support

Device support refers to the kind of devices supported by the solution. Examples may include Windows personal computers, PDA's, routers or Linux workstations. The Jericho Forum requires a solution that can verify any kind of devices as defined in the logical requirements.

Verification model

The verification model refers to the manner in which the solution verifies the endpoints. Examples include continuously, only during pre-admission, post-admission or a combination of these. The Jericho Forum requires a solution that can provide continuous endpoint verification status in addition to the initial verification.

User interaction

User interaction refers to the manner in which the solution communicates with the end-user. The solution should only request user interaction for important tasks. The Jericho Forum requires a solution that positions control of interactions as close to the data and end-user as possible.

Standard adherence

Standard adherence refers to the ability of the solution to operate in heterogeneous environments, whilst maintaining functionality. The Jericho Forum requires a solution that not only adheres to standards, but also uses open standards.

Accuracy

Accuracy refers to the ability of the solution to accurately identify threats and act accordingly. It should be able to negate 0-day threats⁷. The Jericho Forum requires a solution that maintains endpoint security in any environment. Therefore, threats should be able to be timely identified and responded to.

⁷ 0-Day threats are exploits that are released before the vulnerability itself is made public.

Current Implementations

Gartner has established an overview of available NAC solutions in their MarketScope for NAC 2007 Research paper. As of February 2007, 17 enterprises deliver NAC solutions. Of these, Gartner rates⁸ five companies as positive, eight as promising and four as caution. This paper will discuss all solutions rated positive, with the exception of Sophos NAC 3.0, which at this time is still integrating their NAC solution with their Endpoint Security suite. In addition, two promising solutions that appear to match the established requirements will be discussed.

Positive

- Bradford Networks
 - NAC director
- Cisco Systems
 - Cisco NAC appliance/Cisco Clean Access
- StillSecure
 - Safe Access 5.0
- Symantec
 - Network Access Control

Promising

- Juniper Networks
 - Unified Access Control 2.0
- Mirage Networks
 - Endpoint Control

Network Analysis

Although not a dedicated NAC solution, the Lanclope Network Behavior Analysis solution can be used to detect and prevent unusual network traffic. As such, its usefulness within the Jericho Forum model will be discussed.

⁸ Positive: Demonstrates strength in specific areas, but is largely opportunistic

Promising: Shows potential in specific areas; however, initiative or vendor has not fully evolved or matured. Caution: Faces challenges in one or more areas

Bradford Networks - NAC director

Bradford Networks was founded in 1999 as a telecommunications engineering services company. Early projects were driven by a demand for custom capabilities to control devices connecting to academic networks. By 2002 Bradford Networks started on focusing these functions into the Campus Manager product. In January 2007, Bradford Networks announced a new appliance, NAC Director, targeted at the enterprise market. NAC Director intends to provide a comprehensive NAC solution through active enforcement of network usage policies. NAC Director ensures that all devices accessing the network meet required security standards. The solution's endpoint compliance capabilities perform registry-based scans on each network device prior to being placed on the live network. Gartner rated Bradford Networks as "positive" for university environments, but notes that large enterprises that want tighter NAC/mitigation integration should evaluate other providers.

Features

- Persistent and dissolvable agents
- Continuous endpoint posture analysis
- Vulnerability scanning

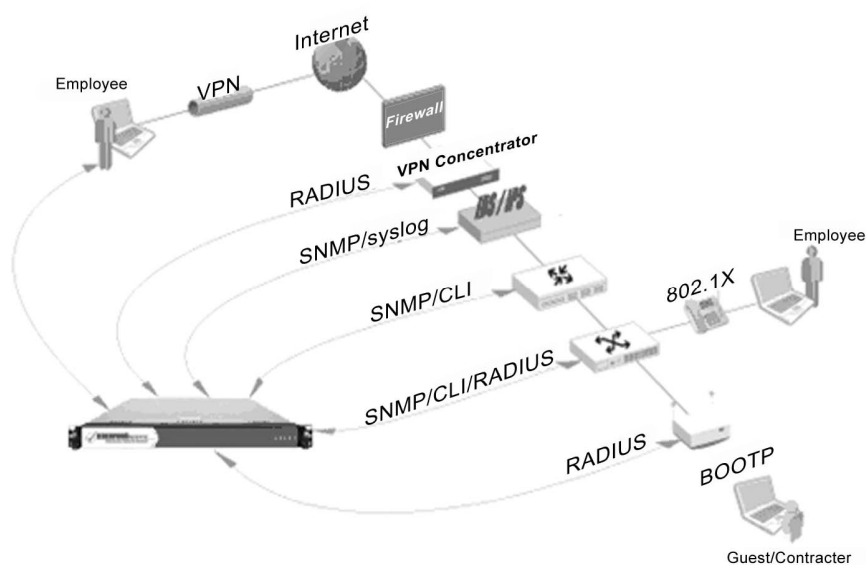
Strengths

- Support
- Out of the box integration with third party vendors
- Well suited for heterogeneous environments

Concerns

- Lack of automated remediation capabilities

How the NAC Director integrates with a network:



Cisco Systems - NAC Appliance & Clean Access Agent

Cisco's network access control solution consists of the Cisco NAC Appliance, a NAC Appliance Manager and optional endpoint agents who provide richer functionality. The Cisco NAC Appliance intends to ensure a secure and clean network environment by analyzing systems attempting to access the network. The system usually installs a small application, known as the Clean Access Agent, on a computer to authenticate the user and verify the software environment. In addition, the NAC appliance can identify whether networked devices such as laptops, IP phones or game consoles are compliant with your network's security policies. Gartner rated the NAC Appliance "positive" for Cisco environments, with the note that competing solutions may prove to be better suited in non-Cisco environments.

Features

- Persistent agents
- Continuous endpoint posture analysis
- Integration with authentication mechanisms
- Vulnerability scanning

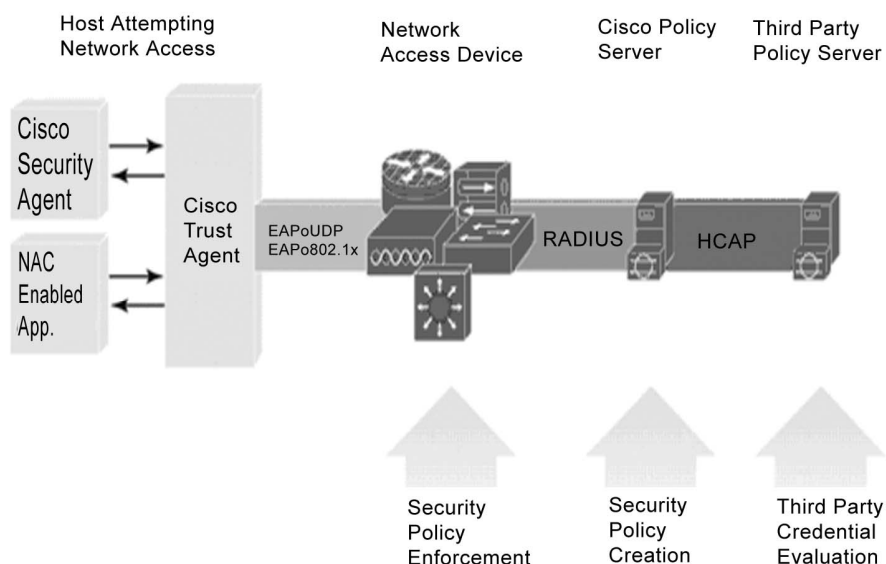
Strengths

- Support
- Automated remediation
- Flexible deployment modes

Concerns

- Limited 802.1X support
- Limited interoperability with non-Cisco equipment

The Cisco Trust Agent Architecture:



StillSecure - Safe Access 5.0

StillSecure was founded in 2000 with the intention to deliver a suite of network infrastructure and security solutions that would improve network security. StillSecure Safe Access intends to protect the network by ensuring that endpoint devices are free from threats and in compliance with security policies before they are allowed on the network. Scans may continue whilst devices remain connected. StillSecure Safe Access offers persistent, dissolvable and agentless testing, compatible with Windows and MacOS X systems. Gartner rated Safe Access 5.0 "positive", believing it will improve its OEM and R&D partners in 2007.

Features

- Persistent and dissolvable agents
- Continuous endpoint posture analysis
- Vulnerability scanning

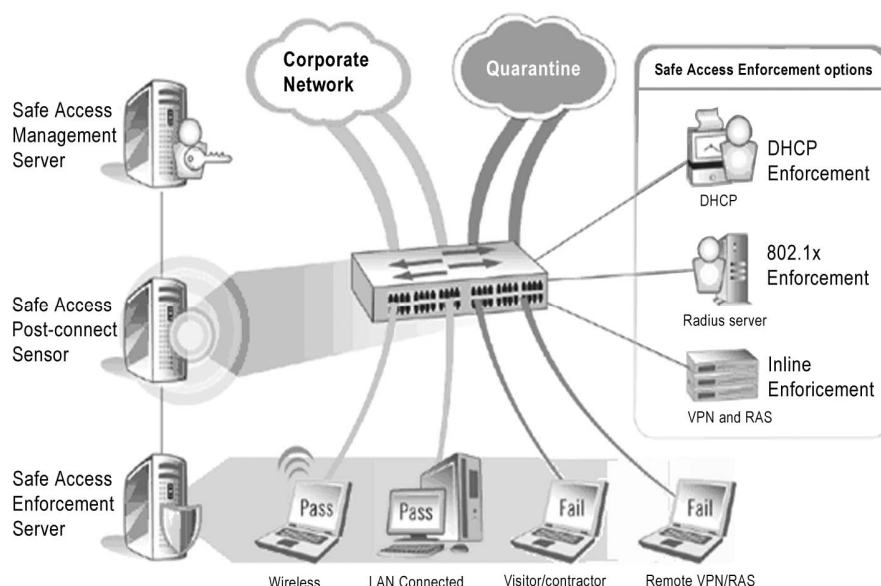
Strengths

- Well suited for heterogeneous environments
- Multiple testing options
- Automated remediation

Concerns

- Aimed at LAN environments
- Relatively few reporting options

The Safe Access infrastructure:



Symantec - Network Access Control

Background

Symantec gained established pre-connect NAC technology when it acquired Sygate and it is working on the integration of the technology with its endpoint security agent. Until this integration is complete, Symantec customers will need to install a separate NAC client in addition to their existing antivirus client. Management of this solution will require separate interfaces. Gartner rated Symantec NAC “positive”, challenging Symantec to overcome its relatively weak foothold in the NAC market.

Features

- Persistent and dissolvable agents
- Continuous endpoint posture analysis
- Vulnerability scanning

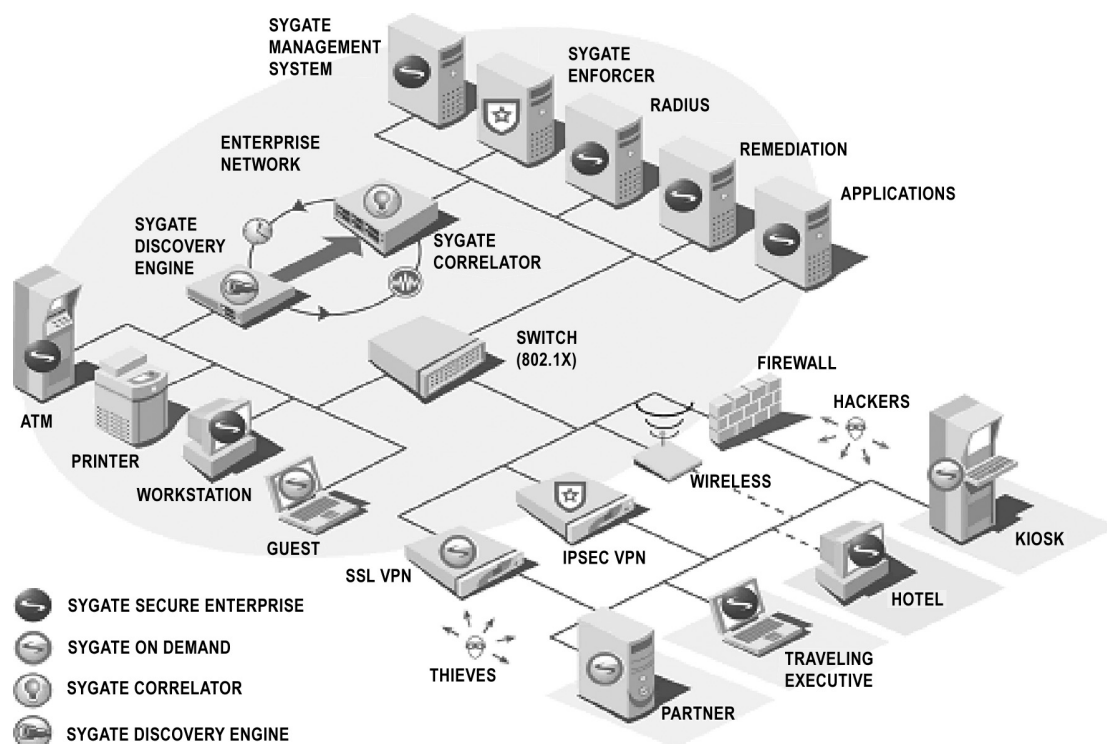
Strengths

- NAC agent also provides self-enforcement.
- Automated remediation
- Self Enforcement
- API integration

Concerns

- No proven LAN implementations
- Additional software required

The Symantec NAC architecture:



Juniper Networks - Unified Access Control 2.0

Founded in 1996, Juniper has become one of the main producers of network related equipment. Juniper has evolved its NAC strategy into Unified Access Control (UAC) 2.0. UAC 2.0 allows Juniper to provide a proprietary solution using Juniper devices for enforcement as well as provide an open solution using 802.1X enforcement. The Juniper Networks UAC solution combines user identity and device security state information with network location information, to create an access control policy for each user. UAC 2.0 can also be provisioned in mixed mode, using 802.1X for network admission control and Layer 3 for resource access control. Gartner rated Juniper as “promising”, suggesting it to invest in reinvigorating UAC as a competitive standard.

Features

- Persistent and dissolvable agents
- Flexible

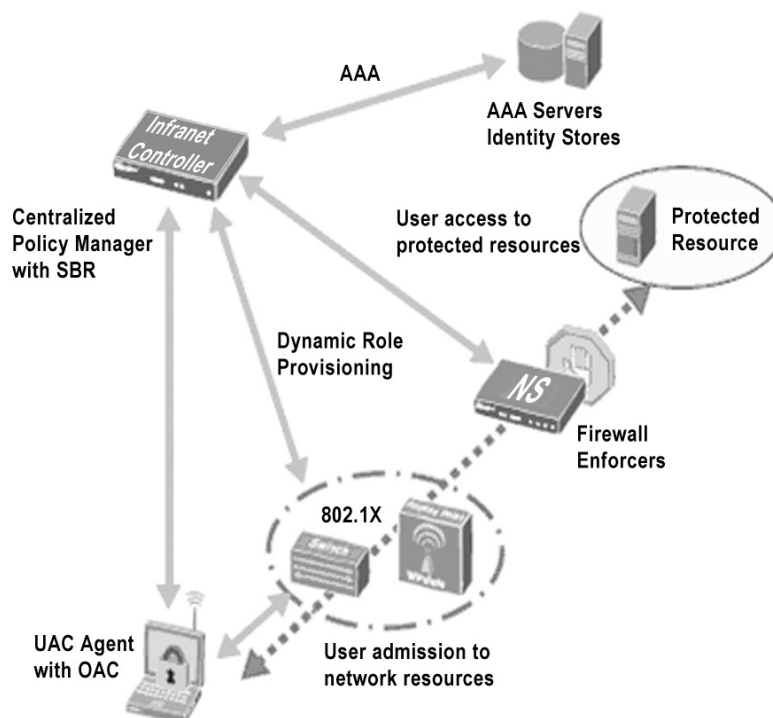
Strengths

- Interoperability with other vendors
- Standard based

Concerns

- Does not provide native patch level verification

The Juniper UAC 2.0 infrastructure:



Unified Access Control Solution v2

Mirage Networks - Endpoint Control

Mirage Networks combines agentless Network Access Control (NAC) with threat prevention and automated policy enforcement. Network intelligence provides a view of endpoint activity and delivers an analysis of network history and usage. Mirage Endpoint Control solutions are network-based appliances. These appliances are deployed out-of-band, thus not requiring extensive changes to existing infrastructures. Endpoint Control provides users with an API that can be used to customize the solution, enabling the solution to be enhanced for specific environments. Gartner rated Mirage Networks as “Promising”, suggesting the development of a persistent agent.

Features

- Dissolvable agent and agentless mode
- Out-of-band deployment with virtually in-line capabilities

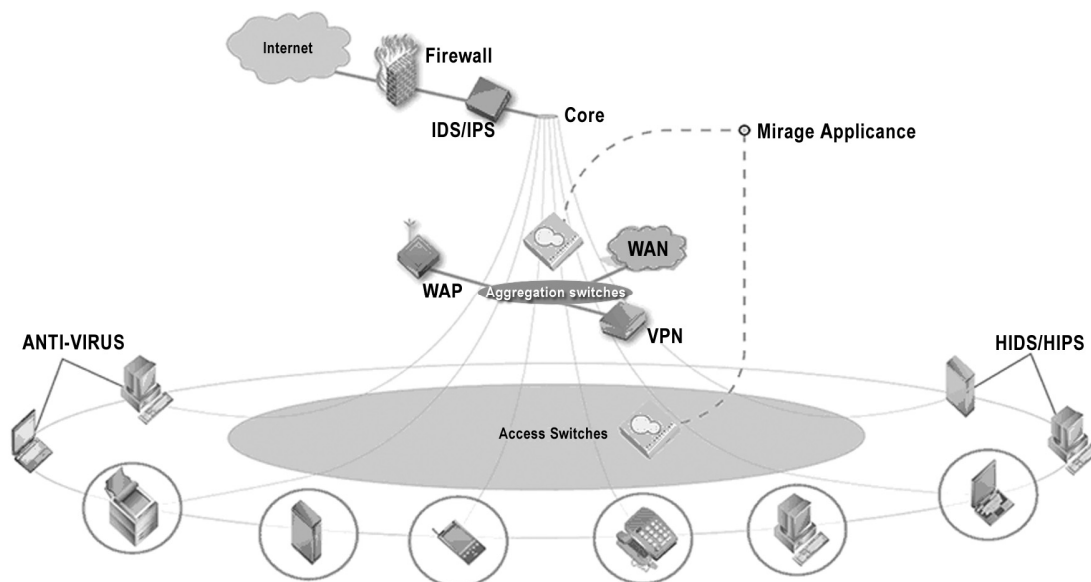
Strengths

- Post-connect monitoring
- Available API

Concerns

- No persistent agent available

The pillars of the Mirage Networks architecture:

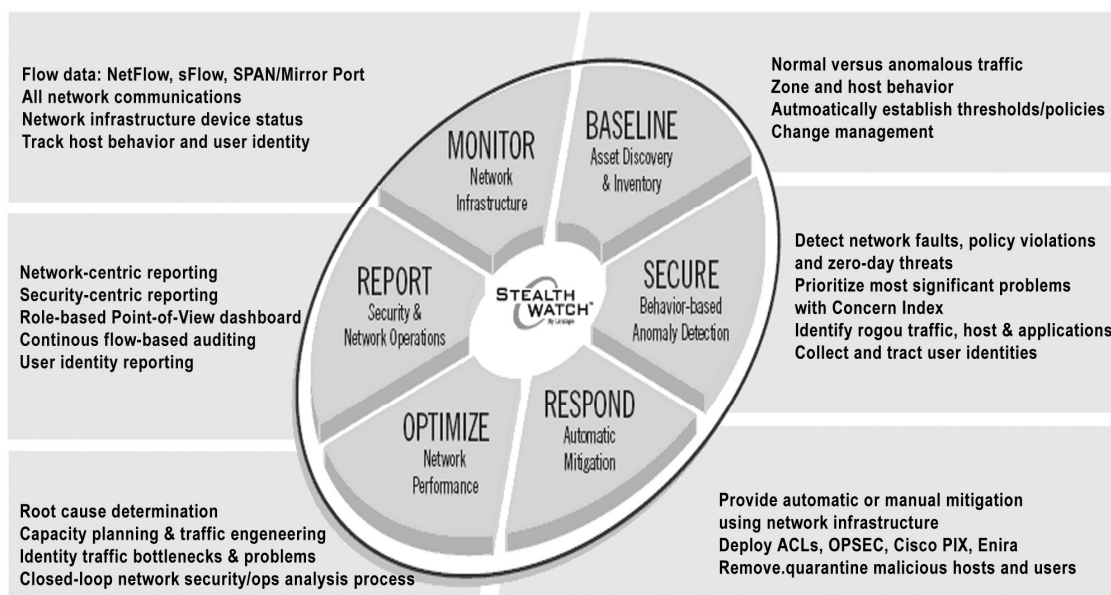


Lancope – Stealthwatch

A network based upon the Jericho Forum vision should not exclusively depend on the AAA-processes to provide and maintain secure network interactions. Methods should be available that can monitor network traffic and make decisions based upon anomalies detected. Traditional Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) commonly use signatures to detect potentially malicious traffic. Jericho Forum based networks use secure and encrypted methods of communication. This means network devices can not read information contained within network traffic, requiring alternative methods of network monitoring. Lancope's Stealthwatch uses baselining to establish an overview of standard network communications. Any serious deviation of this baseline may trigger an alert, either activating an automated response or requiring administrative actions.

Founded in 2000, Lancope has grown to become a major leader in Network Behavior Analysis (NBA) technology, and the manufacturer of the StealthWatch NBA solution. StealthWatch uses native capture or flow records from switches to create an active surveillance system that can adapt to and operate in multiple environments.

This image shows how Lancope's StealthWatch integrates with a network.



Summary of Existing Implementations and Recommendation

Based upon the technical requirements and the features associated with existing solutions, preferred solutions can be determined. The table below establishes an overview of the solutions compared to the technical requirements.

	Bradford NAC Director	Cisco NAC Appliance	StillSecure SafeAccess 5.0	Symantec NAC	Juniper UAC 2.0	Mirage Endpoint Control
Mode of Operation	Persistent Dissolvable	Persistent Dissolvable	Persistent Dissolvable Agentless	Persistent Dissolvable	Persistent Dissolvable	Dissolvable Agentless
Network Integration	Out-of-band	Out-of-band Inline	Out-of-band Inline	Out-of-band Inline	Out-of-band Inline	Out of band
Resources required	Depending on configuration	Lightweight dissolvable agent	Depending on configuration, approx 35K/transaction	Depending on configuration.	Depending on configuration.	Depending on configuration
Device support	MS OS Linux Mac OSX	Any IP based device CCA limited to Windows and Mac OSX	MS OS Mac OSX	MS OS Linux	MS OS Mac OSX Linux Solaris	Any IP based device
Verification model	Continuously	Continuously	Pre-connect Post-connect	Continuously	Continuously	Pre-admission Behavior
User interaction	Limited, Dissolvable agent requires user interaction	Limited, Dissolvable agent requires user interaction	Limited, Dissolvable agent requires user interaction	Limited, Dissolvable agent requires interaction	Limited, Dissolvable agent requires user interaction	Limited, Dissolvable agent requires user interaction
Standard adherence	Uses standard protocols to communicate	Supports open protocols, prefers proprietary protocols	Proprietary engine, uses standard protocols to communicate	Uses standard protocols to communicate	Uses standard protocols to communicate	Uses standard protocols to communicate
Accuracy	Can use multiple applications to provide accurate reporting	Can verify multiple applications to provide reporting	Can use multiple applications to provide accurate reporting	Uses proprietary software for verification	Can use multiple applications to provide accurate reporting	Can use multiple applications to provide accurate reporting

One of the advantages of the Trust broker Architecture is that any solution may be used by any company, it only being required to be able to communicate with a Trust broker. As such, all implementations discussed may eventually be able to be used in a Jericho Forum architecture. One of the aims of this study is to recommend a solution that can be implemented in a prototype network. In order for a solution to be able to function in a prototype network, some key factors must be met:

- The solution should be able to verify any IP based device;
- The solution must adhere to standards;
- The solution must either be able to be customized or the vendor must be willing to customize the application.

There are only two solutions available that can verify any endpoint: the Cisco NAC Appliance and Mirage Networks' Endpoint Control. The solution from Mirage Networks offers an Application Programming Interface (API) that can be used to customize the service. Both products adhere to established standards. In order for the prototype network to gain Endpoint Security verification abilities, it is recommended to implement Mirage Networks' Endpoint Control. In order to provide an alternative solution, Cisco's NAC solution should also be considered for implementation.

5. Authorization in the Jericho Forum Model

Within the AAA-framework, authorization refers to the process of making decisions regarding actions to be allowed or denied, based upon information received from other sources. In a Jericho Forum based network, the authorization process can best be described as being the linking pin between all other processes. Information gathered by other processes serves the single goal of allowing the authorization process to make appropriate decisions. This chapter will firstly discuss process requirements. In order to determine how the authorization process can best serve the Jericho Forum model, several possible architectures will be discussed and compared. Finally, process interactions will be established and a recommendation on how to continue will be made.

Jericho Forum Documentation

Although the Open Group's Jericho Forum has not released any documentation specifically retaining to the Authorization process, the AAA-framework is used to implement the Federated Identity and Trust relationships project directions. In order to determine the official point of view, relevant parts of the Eleven Commandments and the Federated Identity relationship papers will be described.

The Eleven Commandments

The Eleven Commandments is the official document published by the Open Group's Jericho Forum which defines both the areas and the principles that must be observed when planning for a de-perimeterized future. This chapter describes several commandments and the relation between them and the Authorization process. This dissection is based upon version 1.2 of the vision paper, published May 2007.

2. Security mechanisms must be pervasive, simple, scalable and easy to manage.

The authorization process must be able to authorize transactions on a global scale, requiring the least amount of administrative effort possible.

4. Devices and applications must communicate using open, secure protocols.

Communications between the authorization process and other processes must proceed in a secure manner, using protocols that can maintain the confidentiality and integrity of the data.

6. All people, processes and technology must have been declared and transparent levels of trust must exist for any transaction to take place.

All entities involved in an authorization must have been authenticated.

7. Mutual trust assurance levels must be determinable.

All entities involved in an authorization must be able to provide a trust level that can be used within the authorization process.

8. Authentication, authorization and accounting must interoperate with solutions outside your area of control.

The authorization process should be able to use information established by resources outside its own domain.

9. Access to data should be controlled by security attributes on the data itself.

The authorization process should use the security attributes contained on the data itself to determine rights.

Federated Identity vision paper

The Federated Identity vision paper is the official document published by the Open Group's Jericho Forum that describes the Federated Identity problem as envisioned by the Jericho Forum, a recommended solution to this problem and what steps should be taken in the near future. Although this paper does not directly describe the authorization process, it contains several references that may have influence the design of the authorization process. The quotations below are taken from version 1.0 of the vision paper, published December 2006.

Federated Identity

"The majority of user authentication schemes today still use username and password. The burden on users of managing large numbers of username and passwords has led to proposals for Federated Identity systems, where a single set of credentials can be used to authenticate with several organizations, which have agreed to work together as a federation."

Authorization and Federated identity

"The Federated Identity approach has been proposed as a business-to-business service for employees, where one organization manages the user credentials and authorization to systems run by another organization."

Authorization and Data

"In most cases, instead of being centrally stored by a third party, data attributes should be held by the end user. For browser-based applications, a standardized data scheme will enable users to easily transfer data or information between different organizations."

Requirements

Based upon the Eleven Commandments and the Federated Identity vision paper, several requirements can be established. As with the Endpoint Security requirements, these requirements can be organized into the following categories: Trusts, Security, Scope, Manageability, Protocols and Standards. Trust refers to the ability of the solution to handle the additional complexities introduced by the existence of trust related authorization options. The authorization solution should be able to determine whether or not entities are trusted, based upon the existence of trust relations between those entities. Security describes the requirements necessary for the establishment of a secure authorization process. Communications between the authorization process and external processes should be handled in a secure environment. The Scope describes the requirements necessary for the establishment of an authorization process that can handle information obtained from sources outside the local scope. Manageability refers to the necessity of the authorization process to be able to be managed with as few resources as possible, requiring the least amount of administrative input possible for the process to function. The Protocols category refers to the requirements that protocols must meet. As described within the Endpoint Security whitepaper, protocols used should be open and inherently secure. In addition, the implementation of protocols should be transparent in order to be able to replace certain elements of a solution without needing to replace the entire solution, thus future-proofing the network. Finally, the Standards category demands that all protocols used must adhere to open standards. The use of open protocols promotes peer-review and enhances the interoperability between solutions. This interoperability is especially important when considering that authorization implementations should be able to use information obtained from external sources for the decision making process.

Summarized, the authorization process should adhere to the following requirements:

- All entities involved in the authorization process should be authenticated;
- The authorization process should use rights stored on the data itself to determine applicable rights;
- The authorization process must be able to handle authentication information obtained from outside the local domain;
- The authorization process must be robust and require the least amount of administrative effort possible;
- The authorization process should be able to include trust relations;
- The authorization process should use open and secure protocols and standards.


Logical Solution Architectures

There are several architectures available that can be used to implement authorization processes. This chapter will describe the Claims-based, Passive and Active models and, based upon the requirements, will select a model to be used.

Claims-based Authorization

The Claims-based architecture is intricately linked with the user centric identity model, or Identity 2.0, as described by Dick Hardt from Sxip Identity⁹. The concept of Identity 2.0 places the user at the centre of interactions between identity providers and relying parties such as information accessed. The Claims-based architecture is a reflection of the real world where every person is in charge of managing his identity. Imagine a situation where an individual keeps their bank and customer cards in their wallet. When needed, a specific card can be shown to provide only the information required. Within this architecture, information can require certain claims to be provided before access is allowed. This requires the authorization process to be able to determine what claims are required, which claims are received and if the received claims are trusted enough to be allowed to enter the final authorization of requests. Important within the Identity 2.0 philosophy is the need for user consent. Users must be able to determine what claims are required and must be able to select claims they will allow the requester of claims to know. Identity Providers (IP) provide the user with claims that can be presented to Relying Parties (RP).

⁹ OSCON 2005 Keynote, *Identity 2.0*, Dick Hardt



Advantages


- Allows users to control their personal information
- Allows information to strictly control access

Disadvantages

- Requires a complex architecture
- Requires a universal implementation
- Difficult troubleshooting and management
- Requires the development of new solutions

Passive Authorization

A passive authorization architecture relies upon information received from other processes to make decisions regarding interactions. Passive in this context refers to the inability of the architecture to query other processes, such as the authentication process, if additional information is required. Image below represents this process. The user provides the authorization process with credentials (1), which are compared against the Access Control List that determines rights (2). After a match has found, the request is either accepted or denied, and feedback to the user is transmitted (3).



Providing basic and robust functionality, many current implementations are based upon this passive architecture.

Advantages

- Simple implementation
- Relatively easy troubleshooting and management
- Proven technology

Disadvantages

- Lack of flexibility

Active Authorization

Contrary to the passive authorization architecture, an active architecture enables the authorization process to query other processes if additional information is required.

This addition allows the development of new and increasingly flexible and secure authorization processes. Consider the following situation: access is requested to a document rated top secret by the DRM process, from a user logged on to the system using a username and password. According to the access control lists, the user is allowed access to the document; however, a more secure and reliable form of authentication is required. Within an active architecture, the authorization process is able to require the user to provide additional authentication before access is allowed. In addition, data may require a verification of Endpoint status before access is allowed. The Active Authorization process may communicate with the Endpoint Security process to meet data requirements.

Advantages

- Flexible authorization process
- More secure authorization possible
- Allows for integration with the Endpoint Security process

Disadvantages

- Requires an complex architecture
- Difficult troubleshooting and management

Recommended Logical Solution

All proposed architectures offer the basic functionality of authorizing requests. However, as stated within the Aims of the Study, an architecture should be selected that can be implemented using solutions currently available. Although the Claims-based architecture offers several distinct advantages over the other architectures, research into authorization solutions has shown that no currently available solutions can implement this model or offer similar functionalities. Until solutions supporting a Claims-based architecture become available, it is not taken into consideration. Both the Active and Passive model meet the requirements set forth by the Eleven Commandments and the Vision paper, however, the Active Authorization architecture has the additional advantage of supporting the requirements established for the Data Classification project. The Active Authorization architecture is therefore recommended for use in Jericho Forum based networks.

Active Authorization

The Active Authorization architecture enables tight integration with the Authentication and Data Classification processes. This chapter describes the manner in which this integration should function.

Access levels

When a user logs on to the network, an authentication is required. Depending on the methods used for the authentication, an Access Level can be granted. For example, when logging on anonymously, a user is granted Access Level 0. If the user enters a password, Access level 1 is granted. After using a token to authenticate, the Access Level is upgraded to level 2.

Access Levels & File System Authorization

The assigned Access Level can be used for global data authorization. A user granted Access Level 7 is allowed to see data classified at level 7 and below. However, this does not necessarily mean that he is authorized to actually modify the data. Most file systems allow the use of Access Control Lists (ACL's) that are placed on file system objects to allow or deny specific data control features. After being allowed access to data based upon the Access level, file system authorizations need to be used to determine effective user rights.

This authorization process allows for three scenarios to occur:


- Data access with matching Access Levels;
- Data access with insufficient user Access Level;
- Data access without sufficient authorization.

Each of these scenarios will be described in order to establish internal process procedures.

Data Access with matching Access Levels


The image below shows the steps that need to be taken before access to data is allowed. In this scenario, the users' Access Level is sufficient to access the data.

1. The user logs on to a server containing their identity. The Access Level received depends on the methods used to log on. The trust broker uses a token to assign an Access Level to the user;
2. Access to data is requested;
3. The accessed entity verifies the users Access Level;
4. If required by the data, the Endpoint Security process verifies the device the user is connecting from;
5. The entity containing the data verifies that the Access level granted is sufficient to access the requested data. In this case access is granted because the data's Access Level matches the users Access Level. After successful verification, file system authorizations are used to perform data access control.



Data access with insufficient user Access Level


In this scenario, the users' Access Level is insufficient to access the data and additional steps need to be taken.



1. The user logs on to a server containing their identity, gaining an Access Level of 3;
2. Access to data is requested;
3. The accessed entity verifies the users Access Level;
4. The users' Access Level of 3 is insufficient to access the data; he is therefore requested to upgrade his Access Level;
5. The user re-authenticates to the trust broker using biometrics and is granted Access Level 7;
6. Access to data is requested;
7. The accessed entity verifies the users Access Level;
8. The data requires Endpoint verification;
9. The entity containing the data verifies that the authorization level granted is sufficient to access the requested data. In this case access is granted because the user Access Level exceeds the data Access Level. After the successful verification, file system authorizations are used to perform data access control.

Data access without sufficient authorization

In this scenario, no trust relationship exists between the originator of the request and the destination.



1. The user logs on to a server containing their identity. The Access Level received depends on the methods used to log on and the status of the endpoint. The trust broker uses a token to assign an Access Level to the user
2. Access to data is requested
3. The accessed entity verifies the users Access Level
4. It is determined that no trust relationship exists between the user and the data. Therefore, access is denied

Process Interactions

In order to provide a modular framework, interactions with other processes in the Jericho Forum model should be formalized. This enables standardized communications between modules, improving compatibility with other solutions. The first step is to determine the overall process interactions. This requires the establishment of a process flow.

As visualized in image below, the User initiates a request to access the data, transmitting a token containing its credentials and Access level. If needed, the Authorization process may verify the received token. The authorization process queries the data to determine Access Level required and Access Control List options. If demanded by the Data, the Authorization process contacts the Endpoint Security process and requests verification. Information received from the user and the Endpoint Security process is then compared to the requirements from the data. Depending on the results of this comparison, a specific reply is

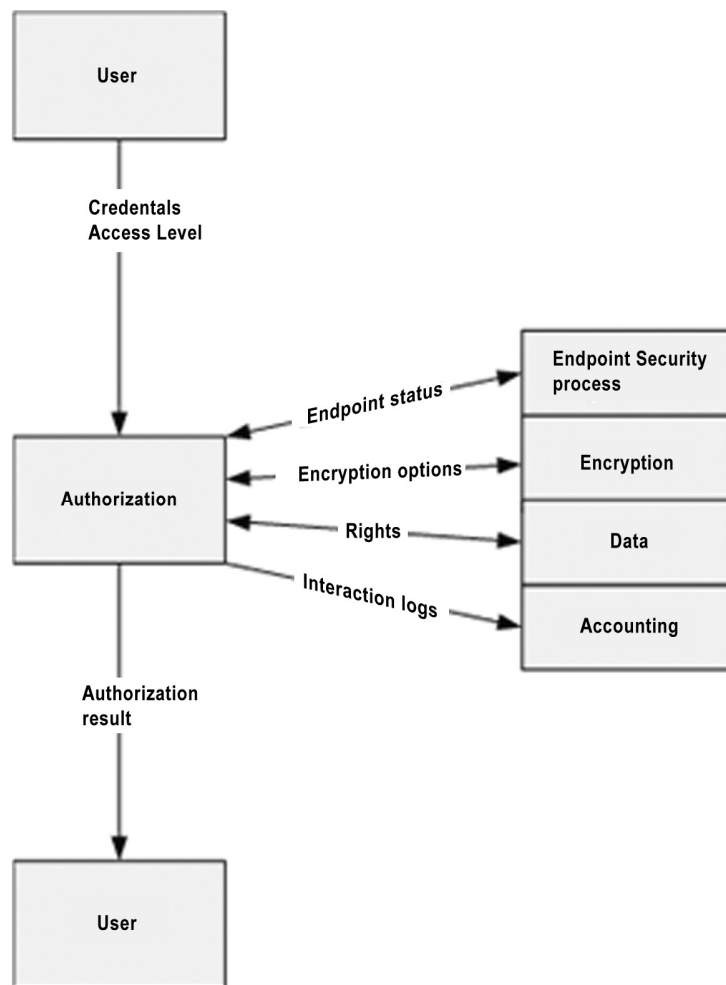
returned to the user. Replies may include the right to access the data, a request for an increase in Access Level, or access may be denied entirely. In addition to these actions, the authorization process interacts with the Encryption process to establish cryptographic protocols and options to be used. All actions performed in this flow must be able to be logged by the Accounting process. Interaction with this process is therefore required. This can be translated to the following inputs and outputs:

Input

- Authentication (Access Level, credentials)
- Data Classification (data rights)
- Encryption (encryption options to be applied)

Output

- Accounting (listing of actions)
- Encryption (request for encryption options)



Technical Solution

Current Implementations and Recommendation

As discussed within the Aims of the Study, the intention of Capgemini's Jericho Forum Research Group is to deliver a model that can be implemented using existing solutions. When considering possible solutions, two aspects must be taken into account: the solution must be able to integrate with the Endpoint Security and Authentication processes and the solution must be able to function in the recommended Active Authorization architecture. No stand-alone and independent Authorization solutions are available at the moment. Rather, vendors incorporate an Authorization process in their solutions, allowing them to provide a complete set of functionalities. Because of this tight integration of Authorization and Authentication processes, the recommended solution for a Jericho Forum based network may require a compromise to be made. The Liberty Alliance¹⁰ has published a listing of Access control solutions that can function in a Federated Identity environment whilst adhering to open standards. Based upon these recommendations, the following solutions were considered:

- HP OpenView Select Access
- Oracle Access Manager
- IBM Tivoli Federated Identity Manager
- Sun Java System Access Manager

In order to determine a suitable implementation, the Authorization aspects of these solutions were compared. Although all solutions provide comprehensive suites of AAA services, none provides an authorization service as described within the requirements and architecture establishment.

Although HP OpenView Select Access does not support the envisioned authorization process, it does provide an API functionality that enables external modules to be implemented. Oracle's Access Manager focuses on protecting resources at the point of network access, rather than allowing resources to protect themselves. In contrast to this method of access management, the Jericho Forum has expressed the need for data being able to protect itself. IBM Tivoli Federated Identity Manager provides a flexible and proven AAA solution, although it focuses on the establishment of a Federated Identity environment, at the cost of providing a flexible authorization solution. Although not providing the envisioned solution, Sun Java System Access Manager does support a wide range of features and offers several customization options. In conclusion, until solutions become available that match the established requirements, two recommendations can be made. Firstly, according to the Jericho Forum, when building a prototype network, the priority should lie with the establishment of a Federated Identity mechanism. It is therefore recommended that a solution being able to cooperate with a Federated Identity driven authentication solution should be implemented. Secondly, in order for future Authorization and Access Manager solutions to incorporate the features described, the dialogues established with these companies during the project should be maintained and improved upon.

¹⁰ See Appendix: Protocols

6. Conclusion and Future Research

Conclusion

Information system requirements are changing as business become more and more open. The Jericho Forum has developed a model that will grant an organization nearly limitless freedom of information system interactions, whilst maintaining the confidentiality, integrity and availability of its data.

Within a Jericho Forum network, the Endpoint Security process enables data to be transmitted to -and through- secure endpoints, enhancing the confidentiality of the data. In order to provide this functionality, several logical and technical requirements were established. The following areas were defined within the requirements: Trusts and relationships, Security, Scope and Scalability, Manageability, Operational behavior, Protocols and Standard Adherence. These requirements established the framework in which the Endpoint Security process has to function. An architecture had to be developed that would allow a global network of Endpoint Security communications to exist. Three architectures were compared and ranked based upon Availability, Scalability and Manageability. Of the architectures compared, the Trust broker model was selected. This model allows local endpoints to direct trust verification queries to a local Trust broker, which becomes responsible for fulfilling these requests by communicating with external Trust brokers. In order to be able to build a prototype network, available solutions had to be compared with the established requirements. Based upon compliance to the requirements and the possibilities of modifying the solution, Mirage Networks' Endpoint Control and Cisco's Network Access Control Appliance were selected. The Authorization process is responsible for making decisions regarding actions to be allowed or denied, based upon information received from other sources. Within the requirements, the following areas were defined: Trusts, Security, Scope, Manageability, Protocols and Standards. Based upon the established requirements, three potential Authorization architectures were discussed: an Active model, a Passive model and a Claims-based model. Because of its ability to communicate with other processes, the Active model was determined to be the preferred architecture. When comparing solutions in order to identify which authorization technologies may be used in a prototype network, it became apparent that currently available solutions are unable to support the model described. Until solutions supporting the requirements become available, it is recommended to use a solution that supports the Federated Identity architecture and can operate with the proposed authentication solution.

In conclusion, it is possible to deploy the Endpoint Security process of a Jericho Forum based network using available solutions. Although no Authorization solutions are available that can implement the recommended model, it is possible to use alternative solutions to construct a prototype network. The next step will be the development of a working prototype in order to enhance organizational awareness of the advantages and benefits associated with a Jericho Forum based network.

Future research

This research paper has established the feasibility of both an Endpoint Security and an Authorization solution within Jericho Forum based networks. In reflecting on the research questions addressed, several interesting avenues of research were observed that may be addressed in the future.

Endpoint Security

At the moment, the Trust broker architecture appears to be the most promising solution to creating secure, global Endpoint Security network. Future research may focus on determining how Trust broker services may interoperate with a claims-based, Identity 2.0 model. In addition, multiple NAC standards are being developed at the moment. For example, Cisco and Microsoft are currently developing a universal standard that should support universal NAC, whilst the Trusted Computing Groups Trusted Network Connect is focusing on incorporating Endpoint Security within a Digital Rights Management (DRM) architecture. Future research could focus on how these standards can eventually become part of the Jericho Forum model.

Authorization

This paper has concluded that at the moment no authorization solutions exist that can implement an authorization process as described. Future research may investigate the possibilities of developing such a solution.

7. References

Literature

1. Anti-Phishing Working Group January Phishing Activity Trends. Anti-Phishing Working Group, 2007
2. C. de Laat, G. Gross, L. Gommans, J. Vollbrecht, D. Spence. RFC 2903 - Generic AAA Architecture. The Internet Society, 2000
3. J. Vollbrecht, P. Calhoun, S. Farrell, L. Gommans, G. Gross, B. de Bruijn, C. de Laat, M. Holdrege, D. Spence. RFC 2904 - AAA Authorization Framework. The Internet Society, 2000
4. P. Windley. Digital Identity-Unmasking Identity Management Architecture O'Reilly 2005
5. Jovanovic, M.A., and al. Scalability Issues in Large Peer-to-Peer Networks - A Case Study of Gnutella. Research report, Univ. Cincinnati, 2001.
6. T.Bowers Rights of Passage. Information Security, 2005
7. Mark Nicolett, Jogn Pescatore, Lawrence Orans, Understanding Benefits of Installed Endpoint Agents for NAC, Gartner 2006
8. Mark Nicolett, Jogn Pescatore, Lawrence Orans, Marketscope for NAC, Gartner 2007
9. Paolo Trunfio, Domenico Tali, Peer-to-Peer Models for Resource Discovery on Grids, ISA 2006
10. Libor Dostalek & Alena Kabelova, *Dns in Action*, Packt Publishing Ltd, 2006

Whitepapers

11. Open Group Jericho Forum Commandments version 1.2 May 2007
12. Open Group Architecture Position Paper Open Group 2006
13. Open Group Protocols Position Paper Open Group 2006
14. Open Group Voice over IP Position Paper Open Group 2006
15. Open Group Wireless Position Paper Open Group 2006
16. Open Group Internet Filtering & Reporting Position Paper Open Group 2006
17. Open Group End Point Security Position Paper Open Group 2006
18. Open Group Enterprise Information Protection & Control Position Paper Open Group 2006
19. Open Group Trust & Co-operation Position Paper Open Group 2006
20. Open Group Federated Identity Position Paper Open Group 2006
21. Open Group Information Access Policy Management Position Paper Open Group 2006
22. HP OpenView Select Access Whitepapers
23. Oracle Access Manager Whitepapers
24. IBM Tivoli Federated Identity Manager Whitepapers
25. Sun Java System Access Manager Whitepapers
26. Bradford Networks NAC director Whitepapers
27. Cisco Systems Cisco NAC appliance/Cisco Clean Access Whitepapers
28. StillSecure Safe Access 5.0 Whitepapers
29. Symantec Network Access Control Whitepapers
30. Juniper Networks Unified Access Control 2.0 Whitepapers
31. Mirage Networks Endpoint Control Whitepapers

Appendix

Protocols and Standards

Several protocols have been developed that allow Federated Identity Systems to become standardized. This appendix describes some of the best known and most widely accepted protocols.

SAML

SAML (Security Assertions Markup Language) is the product of the Organization for the Advancement of Structured Information Standards (OASIS). It is an XML framework that has been designed to exchange authentication and identity attributes between security domains. Several versions of SAML have been developed.

- SAML 1.0 was adopted as an OASIS Standard in November 2002
- SAML 1.1 was ratified as an OASIS Standard in September 2003
- SAML 2.0 became an OASIS Standard in March 2005

Although all implementations of SAML provide the same basic functionality, SAML v2.0 is incompatible with previous versions. Technical solutions may use different versions of SAML, which may cause interoperability problems.

Despite its complexity and limited backwards compatibility, SAML has managed to become the de facto standard in Identity Federation.

Liberty Alliance

The Liberty Alliance was formed in 2001 by 30 members with the intention to establish open standards and best practises for Identity Federation. It has since grown to include nearly 150 organisations. The Liberty Alliance works with other organisations to adopt published standards and to contribute relevant work.

There are 3 important specifications published by the consortium:

- ID-FF
- ID-WSF
- ID-SIS

The Identity Federation Framework (ID-FF) consists of core specifications that make the creation of multivendor identity federation networks possible. The Liberty Alliance realized that a convergence of standards would increase the adoption of Identity Federation. They contributed their set of specifications to OASIS in order to aid the development of SAML 2.0. These specifications enable identity federation and management through features as identity/account linkage, simplified sign on, and simple session management.

The Identity Web Services Framework (ID-WSF) is a framework for discovery and invocation of identity services. After the user has authenticated himself with a identity provider, his assertion can be used by the relying party to discover services this user is eligible for. These specifications provide the framework for building interoperable identity services, permission based attribute sharing, identity service description and discovery, and associated security profiles.

The Identity Services Interface Specifications (ID-SIS) describes how a service supporting identity information of a principal should function. This service provides user and user attribute information. These specifications enable interoperable identity services such as personal identity profile service, alert service, calendar service, wallet service, contacts service, geo-location service and presence service

Open ID

OpenID is an open, decentralized, free framework for user centric digital identity. A user authenticates himself at his identity provider, which can be a blog or a user home page, where the Uniform Resource Identifier (URI) is used as identifier. Principal information and attributes are exchanged between only with user consent.

Open ID may be used in combination with other identity systems such Microsoft Infocards, which would provide more security and support for additional claims that Open ID would be unable to deliver without the implementation of a user agent (e.g. authentication strength, multiple claims and security information).

The advantages of Open ID are its simplicity and lightweight trust model. Its biggest disadvantage is that its usage is limited to web services.

Web Services-* and Microsoft Identity Metasystem

Windows CardSpace (formerly Infocards) consists of client software that enables users to prove their digital identity to online services in a simple, secure and trusted manner. CardSpace can be compared to a wallet where one keeps all his identity cards and presents them to certain authorities when needed. CardSpace provides additional security to the user, shielding him from phishing attacks by authenticating the relying parties.

Microsoft's Identity Metasystem architecture is claimed to use open standards and incorporates multiple protocols that make interoperability between multiple standards possible.

The components of the identity metasystem architecture are as follows:

- Microsoft CardSpace is used as the user agent. There are several open source initiatives which provide OS agnostic user experience. Some agents support multi-factor authentication.
- Identity Providers or Security token services supply users with authentication tokens. At this moment, Kerberos and X.509 security tokens are supported.
- The Relying party is a web service or an application which established requirements for identity and claim assertions.
- Languages which make the conversation between user agent, identity provider and the relying party possible. The languages which form the core of the identity metasystem are called Web Services-* (WS-*). Microsoft and IBM together developed a set of WS-* protocols. IBM describes the WS-* as follows:

"Web services are a loosely-coupled, language-neutral, platform-independent way of linking applications within organizations, across enterprises, and across the Internet. A key benefit of the emerging Web services architecture is the ability to deliver integrated, interoperable solutions -- which makes it critical to ensure the integrity, confidentiality, and overall security of these services."

Web services may be used either together or independently. Each web service solves a particular problem in web services interoperability. Web Services is the foundation of the identity metasystem and allows interoperability between different identity providers and relying parties.

- WS-Policy is a language that describes the security policy of the certain web services: SOAP message security, WS-Trust or WS-SecureConversation. Through this framework a web service may express its security policy and declare how messages are to be secured.
- WS-Trust is a language which allows one security token to be exchanged for another. The specifications define dissimulation and issuance of the security tokens within different security domains.

WS-MetadataExchange is a language that defines how metadata associated with a web service endpoint may be represented as a resource and how this metadata may be retrieved from the web service endpoint.