

Toward a Taxonomy and Costing Method for Security Services

Cynthia Irvine
Naval Postgraduate School
Monterey, CA

Tim Levin
Anteon Corporation
Monterey, CA

Abstract

A wide range of security services may be available to applications in a heterogeneous computer network environment. Resource Management Systems (RMSs) responsible for assigning computing and network resources to tasks need to know the resource-utilization costs associated with the various network security services. In order to understand the range of security services an RMS needs to manage, a preliminary security service taxonomy is defined. The taxonomy is used as a framework for defining the costs associated with network security services.

1 Introduction

Several efforts are underway to develop middleware resource management systems (RMSs) that will logically combine a wide range of network resources to construct a “virtual” computational system [2] [5] [10]. Geographically distributed, heterogeneous resources are expected to be used to support applications with a wide range of computation needs. Large parallelized computations found in fields such as astrophysics [11], aerodynamics, meteorology, etc. will require allocation of perhaps hundreds of individual processes. Multimedia applications, such as voice and video will impose requirements for low jitter, minimal packet losses, and isochronal data rates. Adaptive applications will need to adjust to changing conditions. The RMS in such an environment is responsible for: efficiently scheduling multiple simultaneous tasks onto specific network resources; supporting user requirements for performance and security (viz, QoS); and providing support for tasks to adapt to changing resource availability.

Users or applications submit tasks to the RMS, which schedules the tasks for execution. As part of the process of estimating efficient task schedules, the RMS must balance resource-usage costs against user benefits. Specifically, there might not exist sufficient resources to maximize the benefits to all users. Thus the RMS must quantify the costs associated with the entire range of network services. These costs include bandwidth, task execu-

tion speed, latency, storage, etc. Costing of security services in this context has received little attention. The challenge is to associate costs with the entire range of network security services.

The purpose of this paper is to present a preliminary taxonomy of security services, and to show how this taxonomy can be used as the foundation of a system for supplying security-costing information to an RMS. Section 2 presents our preliminary taxonomy. Section 3 is a sketch for how the structure of the taxonomy might be used to define quality of security service requests to an RMS. Section 4 examines how the cost of using various elements of the taxonomy might be presented to an RMS; and Section 5 is a discussion and conclusion.

2 Taxonomy of security services

Users and applications on the network are presented with various security *services* (e.g., authenticity, confidentiality, integrity, non-repudiation, etc.). A security service may be used to implement one or more security policies (organizational and automated [16]), and is, in turn, implemented by one or more security mechanisms (and of course, a given security mechanism may be used to implement different security services, e.g., “OS access controls,” in Table 1). Some mechanisms provide fixed services, and some are variant.¹ Additionally, the RMS may make choices for the user regarding variant security mechanisms, as part of its schedule formulation or adaptive re-scheduling (see Section 4).

Each security mechanism is associated with a service area, which indicates the general topographical component of the network in which the security or protection is effective. The taxonomy identifies three service areas: end system (e.g., a client or server system), intermediate node (e.g., routers, switches), and network connection (i.e., the “wire” connecting various systems and nodes). Security

1. Variant mechanisms offer the user various “degrees,” or strengths, of security (viz., over and above some minimum requirement). See [9] for details.

mechanisms associated with end systems and intermediate nodes protect resources (e.g., data and programs) that are associated with a node or system; for network connections, we are concerned with mechanisms for protecting information that is physically in transit.

Table 1 provides our preliminary taxonomy. It lists security services, example mechanisms and associated service areas. The service areas are designated: “IN” for Intermediate Node, “NC” for network connection, and “ES” for End System. The Total Subnet (TS) service area identifies mechanisms that cannot be assigned exclusively to either of IN, W, or ES.

2.1 Rationale for the taxonomy

In constructing a taxonomy one wishes it to be both useful and complete. Since a taxonomy is simply an organizational artifice, it must have reason to exist, which is its usefulness. Additionally, the taxonomy fails if it does not account for all of the elements of the classes that it attempts to organize.

We have found this taxonomy to be a useful tool for characterizing the security services and requirements that a RMS might encounter in the network context. As such, it is useful for organizing a quality of security service request (see Section 3) and for presenting costs to a Resource Management System (See Section 4).

As for completeness, we assert preliminarily that the top level is complete. Our taxonomy includes the traditional security categories found in the literature, e.g., Pfleeger [12] (confidentiality, integrity, availability), Ford [4] (authentication, access control, confidentiality, integrity, non-repudiation), and Stallings [15](confidentiality, integrity, availability, authentication, nonrepudiation, access control). (Note that in the latter two examples we find “access control” to be redundant with availability, confidentiality and integrity). Empirically, all of the example mechanisms that we have examined so far have been accounted for in our top level list of security services.

The second level (viz., end system, intermediate node, and network connection) is a simple enough partitioning of the generic network topology that we claim it to be complete through inspection. The list of mechanisms in Table 1 is not intended to be exhaustive, but provides a framework for illustrating the taxonomy.

3 Quality of security service requests

The security service taxonomy can be useful in understanding how security is involved in a Quality of Service request. Security in the Quality of Service context has traditionally implied the general notions of one or more of the

following: confidentiality, authenticity, access control, and integrity [3] [13]. However, there is no reason why a Quality of Security Service request could not include all of the elements from “Security Service” and “Service Area” in Table 1. In other words, we envision a security vector in a fully-functional Quality of Service request to include levels of service for the range of security services and mechanisms that we have identified, where “level of service” can indicate degrees of security with respect to assurance, mechanistic strength, administrative diligence, etc. Thus, a generic QoS request would look something like the following in a BNF-style notation:

QoS Request ::= task_specifier, security_vector,
performance_vector, other_factors

And a security vector would appear as follows:

security_vector ::= security_component [,
security_component]*

security_component ::= security_service, service_area,
level

security_service ::= <services from Table 1>

service_area ::= [ES | IN | NC]

level ::= <mechanism-dependent security-level indicator>

A component may be included in the security vector for each variant security mechanism, i.e., for each mechanism in the network environment that provides to the user a choice of security “level.” For example, a partial security vector might look like this:

data confidentiality, NC, crypto-high (e.g., 128-bit
keys),

authenticity, NC, medium (e.g., public-key signature),

nonrepudiation, ES, high-assurance (e.g., Common Criteria rating EAL7 [1])

Here, for the sake of exposition, the “level” of each security component is somewhat arbitrarily assigned. Establishment of nomenclature and metrics for these levels is the subject of ongoing investigations [7] [18]. Translation mechanisms [6] may be utilized in presenting a high-level Quality of Security Service interface to the user, while managing parameters (such as a suitable translation of “level”) to the underlying detailed security mechanisms.

4 Costing of security services

To motivate the need for security costing information, a specific RMS scheduling mechanism is described. We will

Table 1: Preliminary security service taxonomy

SECURITY SERVICE	SERVICE AREA	EXAMPLE SECURITY MECHANISMS
Data Confidentiality	IN	OS access controls, Cryptographic credentials
	NC	40-bit DES, 128-bit Blowfish
	ES	OS access controls, Cryptographic credentials
Traffic Flow Confidentiality	IN	Active network nodes monitor traffic and inject dummy packets in response to certain triggering conditions.
	NC	Communications use a Virtual Private Network with encapsulated packets
	ES	Traffic padding up to a defined maximum is provided. Beyond that maximum, traffic flow confidentiality cannot be guaranteed.
Data Integrity	IN	OS access controls, Cryptographic credentials
	NC	Cryptographic chaining, integrity sequence numbers, and digital signatures
	ES	OS access controls, Cryptographic credentials
Authenticity	IN	Active network supports internode authentication based on digital signatures
	NC	Data origin authentication, i.e., IP address, digital signatures
	ES	OS identification and authentication mechanism; use of Digital Signature Standard; use of trusted certificate authority
Non-Repudiation	IN	Active network nodes report transactions to secure logging facility
	ES	Digital notary and non-repudiation services
Guarantee of Service, Availability	IN	Active network nodes reserve bandwidth for network administrative traffic. Priority-based scheduling for application traffic
	NC	Bandwidth reservation protocol
	ES	Time-slicing scheduler, FIFO scheduler with preemptive interrupts
Audit and Intrusion Detection	IN	Auditing of network control functions
	TS	Rule-based and profile-based network intrusion detection, intrusion correlation engine to identify intrusions across a group of subnets
Boundary Control	TS	firewall, proxy server, guard

show how this work requires detailed security costing information.

Resource management systems are responsible for efficiently scheduling multiple tasks onto computing and network resources in a distributed, heterogeneous computing environment. RMSs support Quality of Service by scheduling to meet user requirements for performance and security, and by providing support for tasks to adapt to

changing network resource availability.

An RMS schedules tasks for execution in the network in response to requests from users or applications. The task may be submitted with a QoS “specification,” which articulates the user’s desired quality of service, including security services. An RMS currently under investigation, the Management System for Heterogeneous Networks [5], has as its primary goal determination of the best scheduling

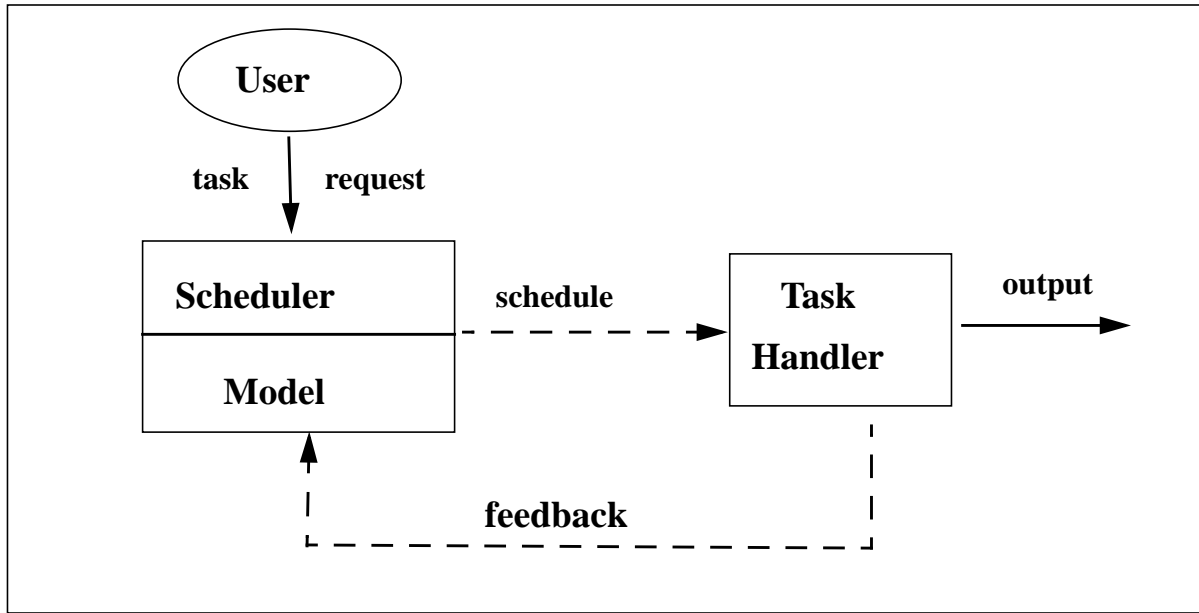


FIGURE 1. Resource Scheduler. The task handler is responsible for realizing the scheduler’s execution plan and provides feedback to the scheduler so that it can dynamically adapt the schedule to evolving resource conditions.

support for many diverse applications, each with its own quality of service requirements, in a distributed, heterogeneous environment. MSHN preserves compatibility with existing security policies, applications and operating systems through its middle-ware role. This is in contrast to network operating systems, which *strictly* control the access to and utilization of resources, and usually require modifications to the OS, application, or security policy.

The MSHN RMS constructs task schedules based on a network infrastructure model. This model includes the resource and security requirements of current and waiting tasks, and the security and availability of network, computing and storage resources. The resulting schedules are provided to task handlers that run the tasks and provide feedback to the scheduler. If the model is inaccurate (e.g., security or resource availability changes), the RMS adjusts its model and potentially reschedules the tasks (see Figure 1).

RMS schedule construction consists of several logical phases, or steps:

1. In the reduction phase, the scheduler finds the *realizable* resource assignments for the task by discarding the possible assignments that will not work according to the model. In addition to resource availability matching (e.g., required service type vs. resource type), security plays a key role. Both the task and the resources are characterized by security requirements: those of the task must be met by a subset of the resources; those of

the resources constrain the task. The task’s security characteristics are compared to the minimum security requirements of the various resources and infrastructure components to determine where the task can run. Additionally, the task’s minimum and maximum security requirements (e.g., reflecting the user’s QoS security specification) are compared to the services available from the resources and infrastructure. The result is a set of resource-assignment “solutions,” where each solution identifies various resources sufficient to run the task.

2. The resource usage costs, including costs for accessing security services, are derived for the various solutions.
3. In the optimization phase, an “optimum” solution is heuristically selected. The criteria for selection is to (attempt to) minimize costs and to maximize the QoS benefit to the users ([8] [9] [17]). More specifically, using realizable resources from the reduction phase, the scheduler attempts to create a schedule to meet QoS requirements for all of its tasks. In order to support as many tasks as possible, the scheduler must meet the typical task scheduling constraints while minimizing resource usage costs.

After step 3, some RMSs may make various network resource reservations. Finally, the task is submitted for execution.

If a particular security mechanism is “fixed” (i.e.,

Table 2: Security cost examples

Security Service	Service Area	Mechanism	Cost Measure
Data Confidentiality	NC	link layer 40-bit DES	Processor clocks per byte [14]
Message Non-Repudiation	ES	remote non-repudiation service	$2n$ bytes per message network bandwidth, plus c clocks per byte
Intrusion Detection	TS	experimental ID system	n Mbytes per second of overall bandwidth, plus m instructions per second, plus b bytes per second storage

always applied) then the overhead for the mechanism is part of the normal cost of running the task and the normal costing mechanism used by the RMS will suffice. For variant security mechanisms, however, the security overhead will vary, depending on the user's QoS request. For example, some task invocations will utilize little, if any, of the variant mechanism; other invocations may utilize the mechanism at an increased level; and, the scheduler may adapt security support (while maintaining any minimum system security policy requirements) in order to schedule the tasks most efficiently. The RMS must calculate how much the use of the security mechanism will increase the cost of the task, according to the specific security "level" requested. For this reason, the RMS must have access to detailed information about the resource cost (as well as the task's requested QoS) for each variant security mechanism. Near-optimal solution selection depends on the accurate estimation of per-task, per-resource, cost of security.

With respect to implementation, the RMS's costing information may be table-driven or algorithm-based, and the cost measurement scale may vary for each mechanism and resource (see Section 4.1 , below).

4.1 Costing example

The security overhead for several security mechanisms is shown in Table 2.

The data confidentiality mechanism is a 40-bit DES encryption mechanism implemented in the link layer. For message non-repudiation, a commercial non-repudiation service mechanism is used. The cost of using this mechanism is a per-message exchange of n bytes with the remote non-repudiation server, and c clocks per message-byte to create the crypto-checksum for the message. The intrusion detection mechanism is shown to use a fixed overhead of the network bandwidth (e.g., for sampling and probing) along with constant processor and storage overhead

Costing information is provided to the scheduler, which will use these data and its current system model to select services, including those for security, that maximize the benefit for the collection of tasks it is serving [8].

5 Discussion and conclusion

A taxonomic framework has been presented for describing security services in terms of broad service categories, network "service areas," and security mechanisms. It has been shown that this taxonomy can be used for different purposes, including the specification of user Quality of Service requests, and the specification of security costs related to network tasks. With respect to Quality of Security Service, we have envisioned that users would be able to specify levels or ranges of desired security service, and as with other QoS parameters, could use these specifications to be able to trade off levels of task performance against requested levels of security.

Continued effort is required to determine the best units for the cost measures. For example, all measures could be unitless and normalized within a common framework. This approach would require a careful description of the semantics of the units with respect to each security service. Alternatively, units can be retained and the components combined into a "vector" to be used by the RMS scheduler.

Additionally, further work is required to expand the enumeration of specific security mechanisms with respect to the described taxonomy.

Acknowledgments

The authors wish to gratefully acknowledge partial support for this research from the DARPA/ITO program.

References

- [1] Evaluation Criteria for IT Security, ISO/IEC CD 15408, Part 3: Security Assurance Requirements, November 1997. (commonly known as the *Common Criteria*)
- [2] Foster, I., and Kesselman, C., Globus: A Metacomputing Infrastructure Toolkit. *International Journal of Supercomputer Applications*, 11(2):115-128, 1997.
- [3] Foster, I, Kesselman, C., Tsudik, G., and Tuecke, S., A

- Security Architecture for Computational Grids, *Proceedings of the Fifth Conference on Computer and Communications Security*, San Francisco, CA, 1998, pp. 83--92.
- [4] Ford, W., Computer Communications Security, Englewood Cliffs, NJ: PTR Prentice Hall, 1994, page 22.
 - [5] Hensgen, D., Kidd, T., St. John, D., Schnaidt, M.C., Siegel, H. J., Braun, T. Maheswaran, M., Ali, S., Kim, J-K., Irvine, C. E., Levin, T., Freund, R., Kussow, M., Godfrey, M., Duman, A., Carff, P., Kidd, S., Prasanna, V. Bhat, P., and Alhusaini, A., "An Overview of the Management System for Heterogeneous Networks (MSHN)," *Proceedings of the 8th Workshop on Heterogeneous Computing Systems (HCW '99)*, San Juan, Puerto Rico, Apr. 1999, pp 184-198.
 - [6] Irvine, C., and Levin, T., A Note on Mapping User-Oriented Security Policies to Complex Mechanisms and Services, Naval Postgraduate School Technical Report, NPS-CS-99-008, June 1999.
 - [7] Juneman, R. R., Novell Certificate Extension Attributes--Novel Security Attributes: Tutorial and Detailed Design. Version 0.998, Novell, Inc. 122 East 1700 St., Provo, UT, August 1997.
 - [8] Kim, Jong-Kook, Hensgen, D., Kidd, T., Siegel, H.J., St. John, D., Irvine, C., Levin, T., Porter, N.W., Prasanna, V., and Freund, R., A QoS Performance Measure Framework for Distributed Heterogeneous Networks, to appear in *Proceedings of the 8th Euromicro Workshop on Parallel and Distributed Processing*, Rhodes, Greece, January 2000.
 - [9] Levin, T., and Irvine C., Quality of Security Service in a Resource Management System Benefit Function, NPS Technical Report, Forthcoming
 - [10] Litzkow, M. Livney, M., and Murtka, M. Condor: A Hunter of Idle Workstations. In *Proceedings of the 8th International Conference on Distributed Computing Systems*, San Jose, CA, June 1998, pp 104-111.
 - [11] Ostriker, J., and Norman. M. L., *Cosmology of the Early Universe Viewed Through the New Infrastructure*. C.A.C.M. 40(11):85-94.
 - [12] Pfleeger, C., Security in Computing, Prentice Hall PTR, Upper Saddle River, NJ, 1997, page 4.
 - [13] Sabata, B. Chatterjee, S., Davis, M., Sydir, J., and Lawrence, T., "Taxonomy for QoS Specifications," *Proceedings of the IEEE Computer Society 3rd International Workshop on Object-oriented Real-time Dependable Systems (WORDS '97)*, Newport Beach, CA, Feb 1997.
 - [14] Schneier, B., Kelsey, J., Whiting, D., Wagner, D., Hall, C., Twofish: A 128-Bit Block Cipher, Counterpane Systems, <http://www.counterpane.com/twofish-paper.html>, June, 1998.
 - [15] Stallings, W., Network and Internetwork Security, Prentice Hall, Englewood Cliffs, NJ, 1995, page 5.
 - [16] Stern, D. F., On the Buzzword "Security Policy", *Proceedings of 1991 IEEE Symposium on Security and Privacy*, Oakland, Ca., May 1991, pp. 219-230.
 - [17] Vendatasubramanian, N. and Nahrstedt, K., "An Integrated Metric for Video QoS." *ACM International Multimedia Conference*, Seattle, Wa., Nov. 1997.
 - [18] Wang, C., and Wulf, W.A., Towards a Framework for Security Measurement, *Proceedings of the Twentieth National Information Systems Security Conference*, Baltimore, MD, October 1997, pp. 522-533.