Securing Critical Infrastructures, Grenoble, October 2004

A GLOBAL FRAMEWORK TO ENHANCE CRITICAL INFRASTRUCTURE PROTECTION

Gwendal Le Grand, Michel Riguidel
GET/Télécom Paris (ENST)- LTCI-UMR 5141 CNRS,
Computer Science and Networks Department, 46 rue Barrault, 75634 Paris Cedex, France

1. Introduction

Our economy and security are increasingly dependent on a spectrum of **Critical Infrastructures** (CI). They are large scale distributed systems that are highly interdependent [3], both physically and in their greater reliance on the information and communication technologies (ICT) infrastructures, which logically introduce vulnerabilities that make them increasingly complex and fragile. Failures, accidents, physical or cyber attacks can provoke major damages which can proliferate by cascading effects and then can severely affect a part or the whole society [4]. Cyber-terrorism is a fundamental threat to modern societies. The international interconnection of networks and the development of information technology has added multiple, dangerous international dimensions to computer crime.

Although over the past few years, ICT have focused on simplification, universality and convergence, it now appears that they are less naively turning toward rich and diversified urbanization. Divergence is inevitable due to diverse technologies and subsidiary controls. In other words, we have a multiplicity of regionalized, autonomous, local and tailor-made solutions to store, communicate, and calculate at lower costs and closer to real use.

The remainder of this paper is organized as follows. In section 2, we present new issues in ICTs. Section 3 focuses on system modelling methodologies. Section 4 proposes an implementation of the models through paradigm that enhance CI protection and tackle intra dependencies (with the policy based management and canonical architecture approach) as well as inter-dependencies (with the virtualization of CI).

2. New issues in ICTs

New paradigms in ICTs

In the ICT domain, new paradigms are emerging that comply with the complex demands of proximity and use, and that encourage the IT and telecom industries to prefer specific solutions that reconcile technology and markets with geography and users. Among these ideas, the concept of **ambient intelligence** points to the filling of geographical space with dynamic digital content (either information or computer programs). The concept of **grids** for intensive computation and the birth of pervasive computing means global and local computation are becoming omnipresent. The planet will be covered with enormous middleware systems which will communicate two by two with variable granularity. Grids are dynamic virtual organizations for performing huge computations, with networks of clustered computers, scalable to enable massive distributed computation. Outsourced computing infrastructures will become semi-public resources and will be separated from their owner-users. This freedom to share computational infrastructures raises several questions such as ethical issues that are not easily solved. Finally, the **urbanization of heterogeneous interconnected networks** proclaims the ubiquity of communications and universal access to telecommunication infrastructures. The planet will soon be covered by these enormous fixed or movable structures, enabling local access to a digital infrastructure that can interoperate with all the other digital structures. Here again, the granularity and size of autonomous networks are very different depending if we consider a PAN, WLAN, WAN, or the Internet.

At the same time, the widespread of **wireless infrastructures** makes it almost impossible to delineate the contours of an information system. Not only has radio enabled building wireless networks, but also using the resurrected distributed computing, we are able, based on a standard resource available in proximity, to weave and configure in space a real and enormously powerful machine performing computation for its own sake. With the standardization of interconnections, it has become impossible to trace connecting wires or interoperating lines between several computers (running sometimes into the millions worldwide). This capability will become a permanent threat, as it will enhance the strength of the individual in relation to the State.

Because of their interdependencies and their increasing reliance on open systems, critical infrastructures constitute an unbounded system where faults may occur and proliferate in a severe way and where security represents a real challenge and requires new methodologies and tools [2]. Modern enterprises adapt quickly with short-decision cycles, fast-reaction loops and just-in-time

procurement cycles. This results in chain reactions and/or hazardous automatic decisions when gaps appear in the behavior of systems and organizations, following inventory shortages, insufficient time, or shortages in logistics with unexpected consequences.

Potential threats to the normal functioning of infrastructures are both natural ("Murphy's Law and Mother Nature") and man-made. Individual outages can be serious enough, but this growing degree of interconnectedness can make possible a whole new scale of synergistic, nonlinear consequences.

Resilience of infospheres

The resilience (security and dependability) of ambient intelligence is a problem that has remained unsolved to this day [8]. To render secure a cyberspace whose airtight boundaries are being dissolved by increasingly wider interconnection, it would be useful to take a step back to gain a systemic, top-down perspective and focus on the essential in order to protect the ambient space with its universally-accepted principles on the one hand, and also assure the security of the infospheres, i.e., the different semantic spaces loitering about in this environment, criss-crossed by passing dataflows, programs or scraps of software that execute themselves from one infosphere into another.

Two types of mechanisms must be established in this landscape. First, Internal mechanisms of each infosphere enabling it to defend and repair itself (self-healing and resilient). Second, organized mechanisms made available and shared by the entire community participating in the ambient intelligence. Of course, such mechanisms will be considered trustworthy by some and untrustworthy by others.

The present-day security models have limitations which are vital to overcome. Given modern information systems, it is important to define security models that are dependent on time, duration, location, geography, context and situational development. To increase their effectiveness and conformity to reality, security policies must be configured according to:

- the needs of the systems' users, taking into account their role and their situation;
- the utilized contents and services, the sensitivity of data and operations;
- the system, its morphology, its function and its behavior.

Therefore, it is essential to define new security models, more fitting to the reality of information systems than the present models, specifying new security policies which are more effective and more appropriate than the present ones, and implementing them on the new digital systems.

It is necessary to chart a transversal infrastructure covering the entire field of these diverse network niches, taking into account the scale factor of these mobile infospheres of widely differing diameters and contents. We should imagine the life cycle of these infospheres of different security which overlap, intersect and overlay each other, and the infrastructure should be able to manage the resolution of any conflicting security policies.

This should not be a monolithic and uniform infrastructure, but it should be adaptable by subsidiarity into each area, by means of ad hoc virtualization of various security paradigms.

The incremental refinement of the security policies should therefore operate in several dimensions:

- in the **symbolic dimension** going from the specification at a high level of virtual abstraction, through the logical and physical plane down to the hardware and software implementation in the equipment;
- in the **geographical dimension**, going from the generic specification to the adjusted instancing of individual networks or sites, anchored in their various technologies;
- and finally, in the **time dimension**, going from the specification of the general security policy inscribed in a charter known to the users, to the real-time proactive and reactive implementation in response to an almost instantaneous event.

The stake, therefore, is to secure the infosphere of the state, referring to the critical infrastructures and their interdependences (the cascading effect of chain catastrophes) for enforcing legislation against cyber-crime. Thus we propose to express the requirements of the system that will determine the choices to model the system.

Requirements of the system

The goals of a security management model are to be able to foresee the development flaws, detect anomalous behaviors to proactively manage the system in order to prevent serious problems, install prevention measures, and reactively control the system by making adjustments in response to changes (that may be sudden as when following an attack) within the system or its environment [1].

Even if it is almost impossible to prevent attacks, it is really important to be able to act quickly within the system to stop a potential proliferation of the problem. Consequently, two correlated works of modeling can be distinguished: one concerning CIs and one concerning security management.

Therefore, the basic requirements of the system are motivated by the following security functional

requirements: prediction and scenario simulation (development, proactive management, etc.), prevention, monitoring (global view, reactive and proactive management, real-time), distributed intelligence and autonomy.

3. System modeling

Security models

An infrastructure is constituted by a set of functional entities (source, transformation, relay, transport, user, cf. Figure 1) and a family of links. The links correspond to flows between functional entities and may be oriented links (eg. water flows) or non oriented links (e.g. information) according to the nature of the supplied resource. Therefore, an infrastructure may be modeled as a non connected graph in which the set of summits (the functional entities) is not convex.

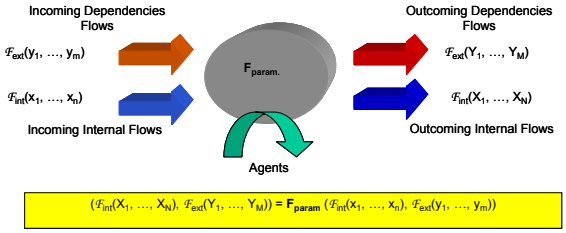

Figure 1. Model of a functional entity

A functional entity may be modeled by a function F_{param}, as represented on Figure 1. This function is a specific function of flow treatment of the entity. It is parameterized by external agents or by information flows. Moreover, the function can correspond to one or several basic roles: source, transformation, routing, transport, consumption. Each entity can have internal modules. A module is an internal sub-entity that fulfills a basic role within the functional entity only for the benefit of the functional entity. In this model, interdependencies and intra-dependencies are modeled using two functions *F*int (family of flows coming directly from another summit of the CI) and *F*ext (family of dependencies flows) that interact with F_{param}. *F*ext may comprise information flows, flows of raw materials which are not handled by the CI, functioning flows (e.g. electricity), risk flows, others.

Each flow may have the following characteristics (oriented/non oriented, material/immaterial, unicast/multicast/broadcast….) and may be of one of the following types: staff (necessary for the correct functioning of F), equipment, hardware, software, … (constituents of F), information, data (supervision, monitoring, operation, …), main resources (directly connected with the resource supplied by the CI), secondary resources (inputs necessary for F), risk flows (geographical interdependency flow, …).

Security policies

To regulate a system's security, a small number of simple and independent laws must be decreed, thereby ensuring that the policy is consistent and known to all users. With the increasing complexity of systems, the laws tend to proliferate in order to cover the ever-increasing number of new situations. The rush into complexity has turned information technology into a novel and irreversible experience. This irreversibility generates the first obstacle in defining security policies. The complexity and abstraction of information systems' architectures has wiped out traditional security policy formulations.

Moreover, the number of canonical computer situations has exploded. Nowadays, it is not uncommon to face computer circumstances that are totally unknown, original and never previously encountered. This new trend is proportionate with the irreversibility of the new information technology. As a result, the number of basic laws is also exploding – should we wish to formulate them exhaustively. If we wanted to write down an effective policy, we would have to formulate a whole procession of laws that would hardly be consistent with the idea that a policy must be simple and known by everyone. Accordingly, we must return to the basic objectives of the policy and formulate the operational security policy using these original objectives, hence the need to revert to policy formulations based on objectives (in a small number) rather than situations (in a combinatorial explosion).

This erosion of the traditional methodologies is also due to the interdependences among the systems themselves. However well a system's laws may be devised, the rules of neighboring infrastructures are generally ignored. The ripple effects can only be blocked if there are common rules regulating the interactions of dataflows between adjacent infrastructures. Finally, the Common Criteria methodology [7], which permits to evaluate ICT security must be renewed in order to take into account the evolution of the past ten years. Security models consider a Target of Evaluation (ToE), immersed alone within a static and inert environment. The CC methodology considers the definition and the specifications of only one single policy for a dedicated target.

Business level vs. technical level

Figure 2 shows the close relationship between the 2 global levels of a CI (business and technical level). For example, when a fault occurs at the technical level, it modifies the degree of criticality of the system [6] and so will affect the business level that can in turn react over the technical level according to its security policy (Dependability). Another fact is that the interdependencies are spread on the 2 levels: logical interdependencies are usually located at the business level, Physical and

Geographic interdependencies are at the technical level, and Cyber interdependencies are at both levels. Therefore, modeling aspects cannot be restricted only to the technical level without considering the relationship with the upper level.

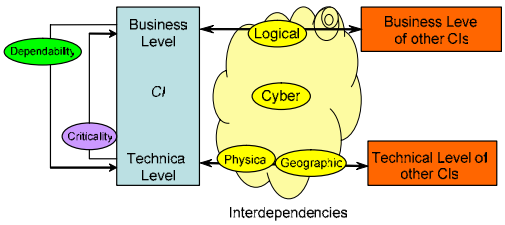

Figure 2. Business/technical level relationships

For modeling, a CI (and even a set of CIs with its pattern of interdependencies) could be characterized as a very large-scale network system, in which a disturbance somewhere in the system can affect everything else in the system. This network, if exposed to a non-trivial disturbance, can no longer respond linearly and either a new equilibrium may not exist or it could be reached only by control actions. Thus, there is a need for a global view of the entire network and a possibility of a quick action over the nodes of this network

System morphologies -- Canonical Architectures

Whatever hierarchical level we look at - from a set of functional entities to the compound of CIs – and even in the organizational schemes of CIs, it is possible to extract a limited set of canonical architectures or patterns or morphologies which have intrinsic properties. These intrinsic properties allow to develop specific, standard and abstract security measures for each canonical architectures depending on such parameters such as the context of the fault [6] and the nature of the flow.

Figure 3 displays a non exhaustive set of relevant canonical architectures. Grey circles and rectangles represent physical or management entities, or compound of these entities that are not necessarily from the same CI. Therefore the canonical architecture theory has two main characteristics. First, Canonical architectures are flow-oriented. Links in the figure can refer to both intra- and inter-dependencies flows. Second, canonical architectures are generic. Entities of a canonical architecture at a certain level can be a compound of canonical architectures at a lower level. This implies a hierarchical modeling with a relevant granularity.

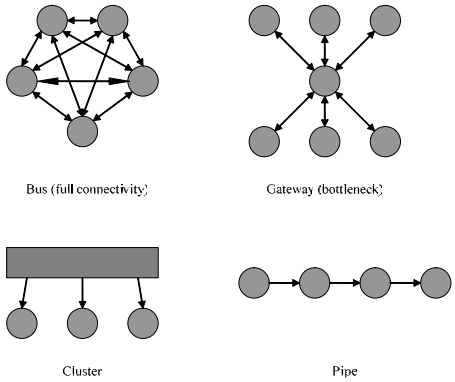

Figure 3. Examples of relevant canonical architectures

By combining specific properties of canonical architectures with a certain amount of parameters such as the nature of the flow, the type and location of the fault, vulnerabilities of existing architectures can be assessed and so security solutions can be developed to design dependable and survivable architectures. Therefore, CI must be modelled as canonical architectures on which we may apply a set of simulations that will allow to modify the morphology of the system in order enhance resistance to faults and attacks.

4. Implementation of the model

Policy Based Management (PBM) Approach

Policy Based Management (PBM) allows a dynamic and global network management. It is *global* since a system is modeled as a state machine in which the union of all local device states gives the global system state. Dynamicity is provided through policies. A network state change provokes a reaction to the event using a bidirectional management. We use the following terminology: *information model*, *data model*, *model mapping*, *policy*, *policy rule*, *policy-based management* (PBM), *configuration management* and *provisioning* defined in [5].

The PBN model (represented on Figure 4) allows to monitor and to reconfigure automatically large numbers of devices to conform to policy parameters. Policies are defined in a high-level language and some mechanism translates them into the various low-level commands that various devices understand. In order to control the devices, policy mechanisms have to be linked to an existing repository of user and resource data. Not only can management tools use directory data to determine policies and locate resources required for the enforcement of those policies, but the tools can also publish information about themselves in the directory. The directory therefore becomes the main point of control on the network, with policy management tools acting as consoles for entering policy definitions and translating them into objects that get published in the directory. The repository's scope integrates all infrastructure services. This helps

eliminate a lot of human mistakes, which can come either from error or from not knowing the relevant policy.

PBN provides a client-server model for policy queries and responses. This scheme is designed to be extensible to all types of policies. A policy server (Policy Decision Point, or PDP) communicates with its clients (Policy Enforcement Points, or PEPs). PEPs send requests, updates, and deletes to a PDP. The PDP then returns decisions to the PEPs. These messages must be authenticated and sent over a secure channel between the PEP and the PDP given the fact that policy servers could represent a powerful means for intruders to create massive disruptions.

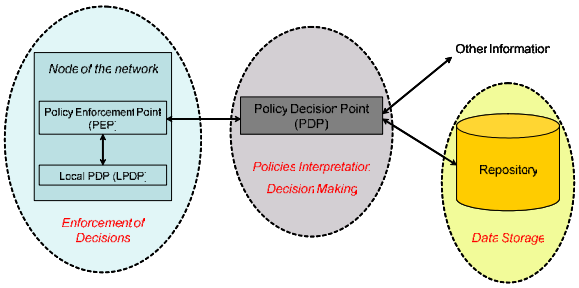

Figure 4. Policy Based Network

The mechanism is stateful: service requests from a client PEP must be retained by the PDP until they are explicitly deleted. The PDP may also push configuration information to the PEP and then remove such state information from the client when it is no longer applicable. In addition, Local Policy Decision Points (LPDPs) let policy/state data and requests be offloaded to subsidiary policy servers closer to the PEPs they control. LPDPs, however, report all policy decision events to the central PDP, which can override them at any time.

Figure 5 represents a hierarchical PBN fitted to the CI environment. We identify two levels of hierarchy in the network. At the high level, the Compound Managing Entity corresponds to a PDP, whereas the CI Managing Entity is a PEP that can be considered as a Compound Managing Agent. At the low level, the Compound Managing Agent is a CI Managing Agent's PDP. The low level itself can be made up of several levels of hierarchy according to the granularity we wish to apply to the model.

Every configuration change, no matter how simple or how sophisticated, has an underlying set of business rules that govern its deployment. Therefore, policies are be adapted to manage such large information systems. A policy rule can be defined as a set of policy conditions and a set of policy actions. PBM can thus be used for critical infrastructure management. PBM is defined as the usage of policy rules to manage one or more entities. Therefore, one or several management entities control the state of the system and objects within the system using policies.

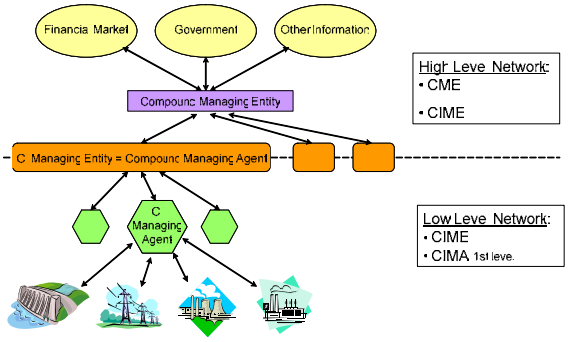

Figure 5. Architecture

We identify two different views for the application of a policy on the technical level: device-specific policies and generic policies. Actually, the second approach seems more appropriate since it is important to separate the modelling of policies from the modelling of device mechanisms (e.g. in the case of a fault in a system, there are standard policies (e.g. standard security measures) only depending on few parameters like the local morphology of the network (cf. existence of a set of **canonical architectures**, the type of faults, …). Therefore we suggest the utilization of device-independent policy models since this is more flexible. The crucial point that has to be considered is the definition of a policy continuum and consitency.

The PBM approach is well adapted to a homogeneous world owned for example by a single operator. In order to have yet a broader vision and transcend the heterogeneity of the CI, we propose the V2V paradigm that allows to model security and resilience of CI at a global level. We propose a general framework to ensure the **security** and **dependability** of the ambient intelligence inhabited by all these infospheres containing their own mobile subjects and objects.

Heterogeneous networks resilience

The V2V (Virtual to Virtual) paradigm aims to inspire trust and ensure the resilience of heterogeneous networks and infrastructures. It is a proposal for creating an end-to-end, adaptive security system across heterogeneous and mobile worlds. Virtualization is a powerful principle used in computer science to conceive of a heterogeneous computerized reality in a different manner by reducing its visible complexity. It hides and absorbs the rough points of heterogeneous entities. It is an artifice, intended to capture several logical abstractions, generally of different types, in order to generate a new organization offering added value. It enables the utilization of several computer or network techniques or standards in order to construct new paradigms based on various technologies. It

enables the creation of logical hooks to get hold of the system components in a different and more efficient manner.

Indeed, virtualization enables the attachment of physical or logical resources that are incompatible, heterogeneous and exploded in order to render their heterogeneity invisible to certain subjects (such as the users), on the one hand, while on the other hand, rendering them more attachable to other subjects, especially to the security hooks that would be able to capture these resources and handle them more efficiently. The management of these resources would evidently be easier and more unified. Virtualization renders space more homogenous and more effective in terms of logical segmentation. It is a concept of continuity, ideal for migrations of technologies.

Security is more serene in limited, static, centralized, homogenous worlds and less comfortable in unlimited, dynamic, distributed, heterogeneous worlds with autonomous, nomadic or mobile entities.

The security of an entity (subject or object) is closely related on the one hand to the trust of this entity toward its environment, and on the other hand to the trust of the environment toward it. The security measures are usually enforced by establishing a catalogue of procedures that permanently test the measure of doubt and/or mistrust of the entity toward its environment and the amount of fear that is expressed by the environment toward this entity.

The environment, the ambient intelligence, the infospheres will set in place procedures to defend themselves from other entities. These other entities, sharing the same environment, must evolve defense mechanisms, either because the environment itself is hostile or, more often, because it is a good idea to protect themselves from the actions of neighboring systems.

V2V paradigm

The objective of the V2V paradigm is to find a framework of solutions to meet the difficult challenge of the security of several distributed systems consisting of open communities of elements communicating within heterogeneous networks and immersed in an ambient intelligence. The entire strength of virtualization lies in the art of choosing the abstractions and the specific mechanisms that will create that added value in terms of effectiveness, reduction of complexity and improved segmentation of the computerized reality. This is still an open issue.

Virtualization enables the transcendence of the present fragmentation of technologies, information systems and networks. It enables us to preserve the ancient architectures and assimilate them into new representations, rendering them compatible with present and future technologies. This virtualization trend should intensify with the advance and accelerated development of technologies.

Security in distributed systems is a field where virtualization seems to have untapped potential. If we want to use virtualization mechanisms, they will be an integral part of security and it will be necessary to make them secure.

The **general security policy** is first of all **defined on a virtual plane** within the entire urbanized computer community, in terms of security objectives (confidentiality, integrity, availability), at a high level of abstraction. By successive refinements, the general policy shall be **projected** on the two traditional **planes – logical and physical**, of information and network systems.

This security policy will be materialized by such projection on the existing architectures and technologies through the PBM management model presented previously.

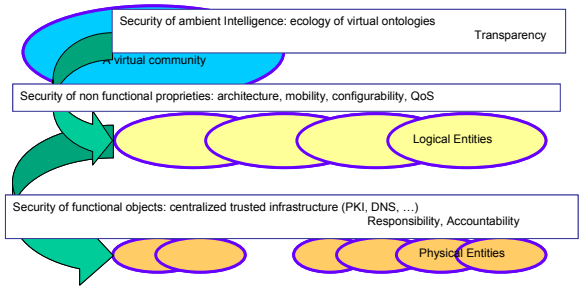

Figure 6: Security of Ambient Intelligence is split at 2 levels; security of non functional properties and security of functional properties.

In the course of these refinements by successive conversions, the security policies will take shape and mutate into **security functions** (identification, authentication, access control, data protection, audit, resource utilization, security management). By **instancing**, the virtual security policy will also become incarnated into the various actors and the various elements of the systems, thanks to security **mechanisms** (encryption, signature, cryptographic protocol, watermarking, filtering, lures, etc.).

5. Conclusion and future works

We defined a modeling architecture that can fit the requirements of any CI. This architecture is based on a hierarchical PBM architecture. An assessment of vulnerabilities of existing infrastructures may be achieved using canonical architectures. Abstract security policies are implemented using a policy based management approach. To transcend heterogeneity of CI and tackle inter dependencies, we propose the V2V paradigm, that helps configure each

security domain (CI), managed using the PBM approach. In fact, the virtualization approach has no interest in implementing security in a homogenous world or in vesting the security concept with specific technology, such as Internet security. On the other hand, this concept becomes extremely meaningful when confronted with the need to assure security of at least two worlds of widely diverging technologies. This new approach also becomes meaningful when proposing to assure security of at least two disjoint security domains within a specific technological universe, such as the Internet, as in the case of split infospheres and the new network applications (P2P, grids, etc.). Virtualization will play here its emancipatory role of deregulation of standards.

Future investigations are planned by us to implement this approach in a Java and UML based application.

References

[1] ACIP project homepage, http://www.iabg.de/acip/

[2] Robert J. Ellison, David A. Fisher, Richard C. Linger, Howard F. Lipson, Thomas A. Longstaff, Nancy R. Mead. *Survivability: Protecting Your Critical Systems*. IEEE Internet Computing, November/December 1999 (Volume 3, No. 6)

[3] Rinaldi, Steven M., James P. Peerenboom, and Terence K. Kelly. *Critical infrastructure interdependencies*. IEEE Control Systems Magazine, December 2001.

[4] David Powell, Yves Deswarte, *On Dependability Concepts with respect to Deliberately Malicious Faults*, STCF 2001, Florianópolis/SC, 5-7 March 2001

[5] John Strassner. *A new paradigm for network management: Business Driven Device Management*. SSGRR 2002s July 29 - August 4, 2002.

[6] David E. Bakken. *Fault Tolerant System Foundations*. CptS/EE 562, Spring 2002.

[7] Common Criteria for Information Technology Security Evaluation, version 2.1, August 1999, http://csrc.nist.gov/cc/

[7] Common Criteria for Information Technology Security Evaluation, version 2.1, August 1999, http://csrc.nist.gov/cc/

[8] IST SEINIT homepage, *http://www.seinit.org*