

# Enhancing Grid Security Using Trusted Virtualization

Hans Löhr<sup>‡</sup> HariGovind V. Ramasamy<sup>†</sup> Ahmad-Reza Sadeghi<sup>‡</sup>  
Stefan Schulz<sup>‡</sup> Matthias Schunter<sup>†</sup> Christian Stübli<sup>‡</sup>

<sup>†</sup>IBM Zürich Research Laboratory  
Rüschlikon, Switzerland

<sup>‡</sup>Horst-Görtz-Institute for IT-Security  
Ruhr-University Bochum, Germany

{hvr, mts}@zurich.ibm.com {loehr, sadeghi, schulz, stueble}@crypto.rub.de

## Abstract

Grid applications have increasingly sophisticated functional and security requirements. However, current techniques mostly protect only the resource provider from attacks by the user, while leaving the user comparatively dependent on the well-behavior of the resource provider.

In this paper, we take the first steps towards addressing the trust asymmetry by using a combination of trusted computing and virtualization technologies. We present the key components for a trustworthy Grid architecture and propose an implementation. By providing multilateral security, i.e., security for both the Grid user and the Grid provider, our architecture increases the confidence that can be placed on the correctness of a Grid computation and on the protection of user-provided assets. In order to maintain important scalability and performance aspects, our proposal aims to minimize overhead. Towards this end, we propose a scalable offline attestation protocol, which allows selection of partners in the Grid with minimal overhead.

**Keywords:** Trusted Computing, Grid, Scalability, Attestation

## I. INTRODUCTION

### A. Background

Grid Computing has been very successful in enabling massive computing efforts to take place, but has hitherto been dominated by ‘big science.’ These projects are usually in the scientific or academic domain (such as SETI@HOME or distributed.net) and, while important, they usually have less stringent security requirements than commercial IT systems.

Currently, security is built into Grid toolkits (e.g. the Globus toolkit [11]) used at the *provider* sites (parties that offer resources for use in the Grid). Secure channels, authentication [6], unsupervised login, delegation, and resource usage [12] are all handled by the toolkit. These mechanisms usually do not concern themselves with protecting the Grid user (the person or entity wishing to utilize resources). The user is forced to trust the provider, often without possibility to verify if that trust is justified. However, the user is not usually regarded as trustworthy in much of the current literature on Grid security (e.g., [17], [37]). This trust asymmetry could potentially lead to a situation in which the Grid provider causes large damages to the user with little risk of detection or penalty. An attacker might publish confidential data, or sabotage the entire computation by providing false results. These problems are most evident in computational Grids, especially in mobile code [14] scenarios. Other Grids, such as storage or sensor Grids, may also suffer from the negative consequences of this trust asymmetry, perhaps to a lesser extent. Due to this problem, companies are hesitant to utilize available Grid resources for critical tasks.

Given this state of affairs, [20] have proposed to utilize the emerging *Trusted Computing* (TC) technology for the Grid, with [34] more closely examining which scenarios require TC techniques. Trusted Computing can be used to enforce *multilateral security*, which means fulfilling the security objectives of all involved parties. The importance of enforcing multilateral security requirements in the Grid has also been stressed by other works, such as [21].

In this paper, we consider “trust” to be the opposite of enforcement. Thus, a trusted component is a component whose well-behavior cannot be enforced by another component and, therefore, has the

ability to violate security policy. This inherently negative view of trust contrasts with the notion put forward in other Grid-related works, such as [5], [35], [36], which view trust as a positive, reputation-based property. It also does not conform to the definition of the Trusted Computing Group (TCG), which considers trust to be a guarantee of the well-behavior of the system. We call such a system a “trustworthy” system.

A large percentage of the platforms used in large-scale Grids are built using general-purpose hardware and software. The TCG has proposed a Trusted Platform Module (TPM) which is expected to soon be incorporated into most newly-shipped general purpose computers. This module usually takes the form of a dedicated hardware chip. The TPM chip is tamper-evident (and ideally, tamper-resistant) hardware that provides cryptographic primitives, measurement facilities, and a globally unique identity (see also [19]). For verification purposes, remote party can query the TPMs measurement of the Trusted Computing Base (TCB) by means of *attestation*. This mechanism, proposed by the TCG, allows for (remote) verification of the status of a platform’s TCB.

One approach to securing computing systems which process potentially malicious code (such as in many number-crunching Grid applications) is providing a virtualized environment. This technique is widely used for providing “V-Servers,” i.e., servers which feature several virtual machines. Each of these virtual machines may be rented to one or several users. While users have full control over the virtual environment, they cannot cause damage outside that environment, except possibly through attempts at resource monopolization, for example, by “fork bombing.”

While virtualization offers abstraction from physical hardware and some control over process interaction, problems remain to be solved. For example, in the x86 architecture, direct memory access (DMA) devices can access arbitrary physical memory locations. However, hardware innovations such as LaGrande [2] and Pacifica [1] aim to address these problems and could eventually lead to secure isolation among virtual machines. Alternatively, direct access to DMA devices could be restricted to the virtualization layer itself. Virtualization technology can be leveraged for building a trustworthy Grid environment, especially since several works, such as [30], have already begun to consider architectures which feature policy enforcement in the virtualization framework.

A trustworthy Grid environment that enforces multilateral security would have a number of benefits. Even sensitive computations could be performed on untrusted hosts. Community computation could become an attractive mode of conducting computationally-intensive tasks. Most personal computers used today possess computing abilities in excess of what is required for casual or office use. These resources could be leveraged to run Grid jobs in parallel to the users’ normal workflow and provide the computational power necessary for next-generation modeling and simulation jobs, without costly investments into new infrastructure. Enterprises could utilize the already-present office machines more fully, resulting in earlier return on their investment.

## B. Our Contribution

Currently, the trust asymmetry explained earlier hinders adoption of Grid computing for commercial scenarios. Current solutions have so far failed to convincingly solve this problem. Enforcement of user-provided policies have always relied on the good will of the provider. This shortcoming was alleviated through reputation systems or deterrents by liability law, but such solutions cannot cope with highly dynamic, and largely pseudonymous scenarios.

We propose a realistic solution that actually *enforces* multilateral security in a Grid scenario, and lay the groundwork for the design of such an architecture. Using isolation between virtual machines combined with a trusted base system (“Trusted Virtualization”), our design is able to protect confidentiality and integrity in a multilateral fashion, thus allowing outsourcing computations to untrusted or even anonymous parties without compromising security. We feel our compartmented security design offers a stronger level of protection than many current techniques can provide.

Based on this architecture, we propose a job submission protocol that is based on Offline Attestation. This protocol allows to verify that a previously-selected partner is in a trusted state, with little overhead and improved resistance to attack. This protocol also guarantees *transitive* trust relations in case that partner performs further delegations.

## II. SYSTEM MODEL AND SECURITY REQUIREMENTS

### A. Notation and Definition

In this work, we use the following symbols:  $\mathcal{P}$  and  $\mathcal{U}$  denote the set of providers and users currently registered in the Grid, with  $P$  being a single provider, and  $U$  a single user.  $K$  is an asymmetric cryptographic key, with  $s_K$  being its private and  $p_K$  its public part.  $\sigma$  is the actual state of a machine at any given time, a summary of all trusted code that either has been, or is currently executed on the machine.  $\sigma'$  is its reported state (these two may differ). Both  $\sigma$  and  $\sigma'$  can be encoded as `config`, a short representation of the state (e.g., a hash value) as determined by the measurement facility of the machine, e.g., the TPM.  $J$  is a Grid job. The format of  $J$  is described in more detail in Section II-C.  $SP_X$  denotes the policy of some entity  $X$ .  $\mathcal{C}$  is a set of credentials of some kind, without specifying the specific type. A specific aspect of a policy is the *good* set, which contains the  $\sigma$  values of all states considered to be trusted by that policy. For a finite number of trusted configurations, this can be expressed as an enumeration, but more complex scenarios likely require a more complex encoding.  $\text{enc}_{p_K}(X)$  denotes a piece of data  $X$  encrypted with a public key  $p_K$ .  $\text{sign}_{s_K}(X)$  denotes an item of data  $X$  which has been digitally signed by a private key  $s_K$ .

### B. Usage Scenario

First, we give an usage example of our grid architecture with offline attestation (shown in figure 1):

A node (such as F), upon joining the Grid, generates and publishes an attestation token ( $\tau_F$ ), which can be used by potential partners to obtain assurance about the node's trustworthiness. Grid users (such as C) retrieve attestation tokens from different grid nodes and select a token indicating a configuration they are willing to trust. The selection decision is made offline, and incurs negligible overhead on the part of the user. Once an acceptable provider is found, the users can submit jobs that can only be read by the selected node, in the trusted configuration. If the node has changed to another configuration, communication will fail.

The main advantage of this approach is that the creation of the attestation tokens is decoupled from the process of job submission, while still providing freshness. Additionally, these tokens are transferable and their correct creation is verifiable without interacting with their creators.

### C. Our System Model

We consider an abstract version of the Grid, which features only  $\mathcal{P}$  and  $\mathcal{U}$ . Each  $U \in \mathcal{U}$  can attempt to access any  $P \in \mathcal{P}$ . Each participant in the Grid is considered to be a partner-and-adversary, which potentially wishes to harm other participants, but also provides services. Such a partner-and-adversary can only be depended upon to execute a given task correctly if it can prove its inability to cause damage (break a partner's security policy). We consider delegation to be modeled as one participant being part of both  $\mathcal{P}$  and  $\mathcal{U}$  (see figure 1). A machine  $m$  is a single physical host. It can host one or more logical participants of either role, and offers a number of interfaces, both externally and between its internal components. Every participant has its own, distinct policy. Each component of  $m$  is an independent actor offering some interface(s) to other components, and usually utilizing interfaces offered by other components.  $\mathcal{P}$  and  $\mathcal{U}$  need not be static, but can instead dynamically grow and shrink as new resources are added to the grid virtual organization (VO), and some participants leave the VO. However, such adjustments (joining and leaving) are not the focus of this paper.

For our purposes, a job image is a tuple  $J = (\text{data}, \mathcal{C}, SP_U)$ . `data` may be an invocation to some predefined interface, or carry executable code. While certainly possessing a quality different from input data, for security purposes, a piece of executable code poses the same requirements and can be protected utilizing the same techniques as any other piece of data. Therefore, we do not distinguish between "code", and "data", but instead consider executable content to be a special kind of data.

$\mathcal{C}$  represents the credentials of  $U$ , which may be needed to gain access to  $P$ . The user also passes a policy  $SP_U$  as part of its invocation, which specifies constraints to be upheld for that particular job. The job, once scheduled, can communicate directly with  $U$  (subject to the  $SP_P$ ).

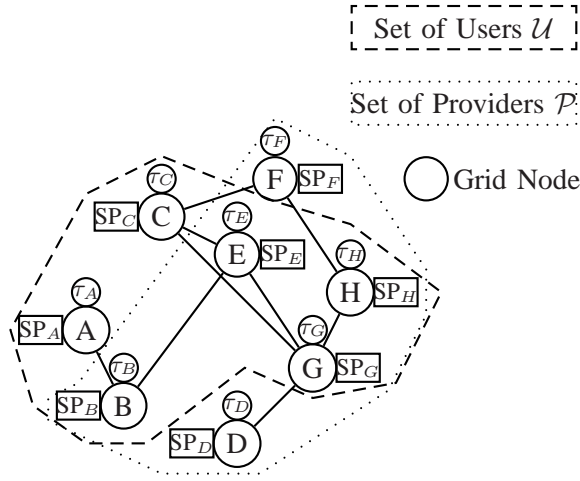


Fig. 1. An example Grid

A machine  $m$  always has exactly one state  $\sigma$ , since  $\sigma$  describes the status of the virtualization layer itself, instead of any particular VM. This state comprises all code running as part of the TCB. TCB components are critical to the correct functioning of the system and need to be trusted. Adding, removing, or modifying such a component changes  $\sigma$ , but  $\sigma$  will not change due to “user actions,” such as installing application software, browsing the web, or executing a Grid job. Furthermore, the system will not allow any party (not even system administrators) to alter the TCB without changing  $\sigma$ .

#### D. Requirements

While most policies describe high level requirements on the use and dissemination of data, such requirements can be reduced to simpler building blocks. Essentially, three main types of requirements exist: confidentiality, integrity and availability. These requirements can be applied to data and combined, and form the basis of all more complex policies. In the context of the Grid, interoperability is often an important functional requirement, both between machines, and different Grid environments. We will briefly discuss interoperability in Section VI.

Availability has been considered in other works, and runs orthogonal to our research. We will focus on providing multilateral protection of confidentiality and integrity in this paper.

##### a) Integrity

Integrity requirements mean protecting some item of data from unauthorized modification. These requirements usually are satisfied only by detecting violations of integrity, instead of actively preventing them. The user  $U$  might be motivated to alter aspects of the provider  $P$ , for example to elevate its privilege level. Similarly,  $P$  might wish to modify the job  $J$ , either to manipulate the computation, or to modify program behavior. Either side needs to be protected from attacks of the other.

##### b) Confidentiality

Confidentiality requirements represent the need to prevent some parties from accessing an item of data. Protected data could, in principle, be freely modified and even destroyed, as long as the attacker is not able to deduce the meaning of that data.  $P$  needs to protect its secret information from  $U$ . A malicious job could collect secrets, such as signature keys, and forward them to  $U$  if confidentiality was not enforced. However,  $U$  might utilize confidential data as part of  $J$ , and demand this data not to be disclosed to any party other than  $J$ 's execution environment.

### III. A TRUSTED GRID ARCHITECTURE

In this Section, we present an architecture for enhancing security and trust properties in Grid computing. We call our architecture the Trusted Grid Architecture (TGA). Figure 2 shows the abstract building blocks of the TGA. The hardware platform provides a TPM and untrusted storage. The trusted

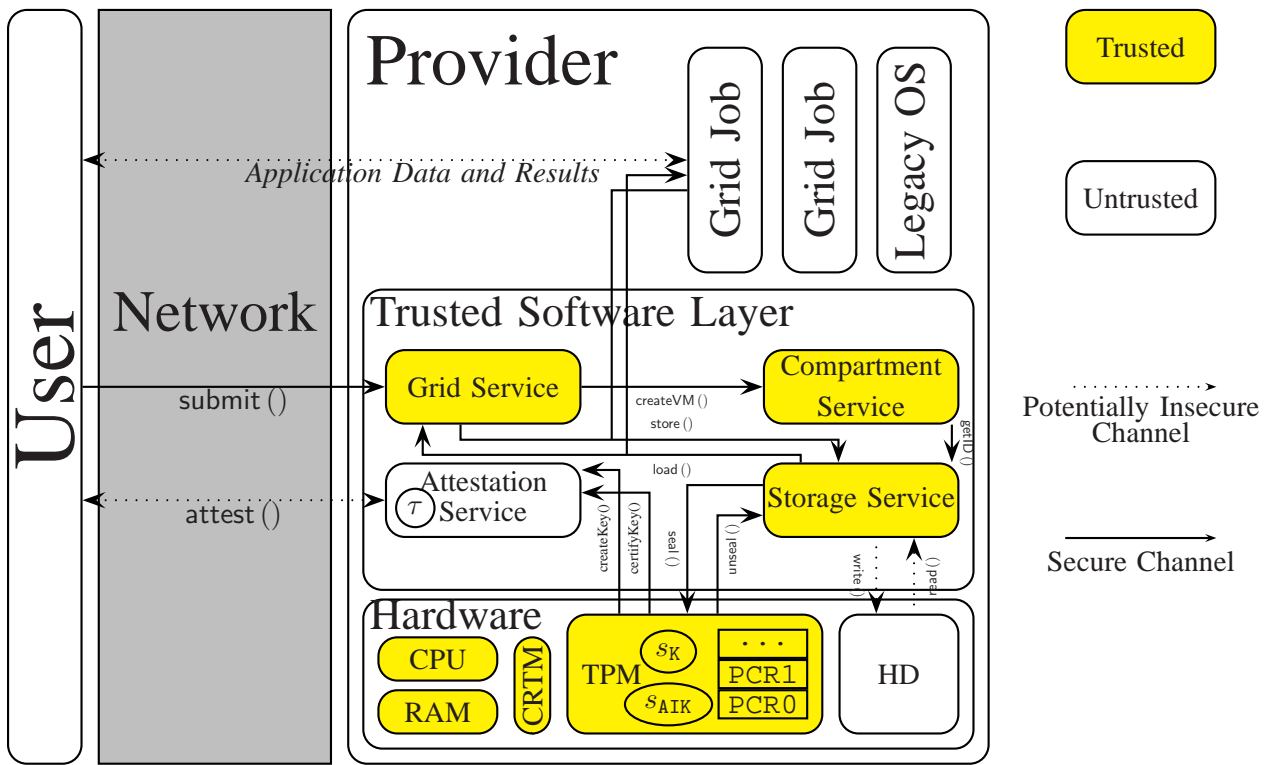


Fig. 2. The Trusted Grid Architecture, and its Components

software layer consists of the attestation, Grid management, compartment management, and storage management components, as well as a Policy Enforcer service. The trusted software layer provides both security functionalities, and virtualization of the hardware (that is, the ability to run several virtual machines in parallel).

The Policy Enforcer Service is not examined here, we assume such a service exists, and is part of the TCB. The Policy Enforcer often depends on the actual policy details to be enforced, as well as the details of the concrete implementation of the otherwise abstract components. Other works, such as [23], [30] and [8] examine some necessary properties of policy engines.

The TCB needs to be trusted by applications utilizing it. It must be resistant to attacks. If the TCB is vulnerable to attack, and code can be executed with full privilege, this potentially implies a full system compromise. Such a full compromise need not even be communicated to the TPM, meaning that the measured state does not change from the trusted values. TPM measurements therefore only guarantee load-time integrity of a TCB, and do *not* rule out a later compromise. Only maintaining load-time integrity places a special focus on preventing TCB compromises. Proper design of a minimal set of trusted services can help to achieve a TCB with the highest possible resistance to attacks.

Additional guarantees about runtime behavior and state may be provided by some dedicated service within the TCB, possibly as an extension to our attestation service. Such techniques are beyond the scope of our current discussion, however. Other works, such as [15] strive to provide such guarantees by monitoring the runtime behavior and state of programs.

Of all the components mentioned, only the TPM and CRTM are trusted at “face-value” – by virtue of its fabrication adhering to TCG specifications; trust in the rest of the TCB derives from the TPM’s *root of trust* only after measurements confirm their integrity. The entire TCB is subject to this measurement, but no untrusted application (such as Grid jobs, or Legacy OS images) is measured, since they cannot violate the security policy (as the appropriate limitations are enforced on them). In case the state  $\sigma$  of the TCB ever changes after boot, it is the TCB’s responsibility to stop any Grid jobs currently executing *before* changing into the new, possibly untrusted, state.

## A. Hardware

Besides the trusted CPU and RAM, two other pieces of hardware play a central role in our proposal: The TPM chip, and the harddisk. The TPM chip naturally is vital, since without it, no integrated trusted third party would be present. It functions as a root of trust, which means we trust its manufacturer to correctly implement the TPM functions. We further assume that the cryptographic primitives used are secure. Each TPM possesses a number of platform configuration registers (PCRs), at least 16 as of version 1.1b of the specification. The PCRs are not arbitrarily writable, but instead can only be updated by means of the *extend* operation. A trusted boot process extends measurements of the platforms TCB, starting from the CRTM. The values of the PCRs thus represent all trusted software executed on the platform, and a set of values can be uniquely tied to a particular configuration.

While the TPM supports a large number of operations, we will utilize four in the course of this article. These are: secure key generation, measurement, certification and sealing.

The TPM features a hardware random number generator, and implements generation of RSA key pairs  $K = (p_K, s_K)$ . For these key pairs, usage limitations can be defined. One possible limitation is sealing, which marks the private key as not being migratable, and usable only given the same PCR values as where present when it was being created. It is possible to obtain a certificate stating which usage conditions apply to a key pair (as represented by its public key  $p_K$ ) from the TPM, signed by one of its *Attestation Identity Keys* (AIKs). AIKs are special signature keys created by the TPM. The private key of an AIK cannot be extracted from the TPM, i.e., it is non-migratable, and it cannot be used to certify migratable keys. AIKs can be certified by a CA, or it can be proved to be a valid AIK anonymously by means of Direct Anonymous Attestation (DAA). Such a certificate or proof is denoted as  $\text{cert}(AIK)$ .

If the TPM is trusted, its measurements can be used to justify extending that trust to other components. As previously explained, it can measure (hash) data into its PCRs. In addition, it can report these measurements to other parties. The values of the PCRs can be signed with an AIK for this purpose. Since the AIK can only be used in a small set of TPM operations, none of which allows for the signing of arbitrary data, a valid AIK signature guarantees that the TPM generated the signed structure. For our purposes, we use signed `KeyInfo` structures that are considered as certificates, although they do not conform to common certificate standards, such as X.509. A `KeyInfo` structure of a sealed key includes the selection of PCRs that were used for sealing, as well as their values at the time of key generation and the values of the selected PCRs that are needed to utilize the sealed key. Also, the `KeyInfo` structure indicates whether a key is migratable or not. We use an AIK to sign such a structure with the `CertifyKey` operation of the TPM and denote the resulting certificate by  $\text{cert}_{AIK}(p_K)$ .

These restricted keys enable *sealing*. Data sealed to a certain configuration of the system is encrypted with a public key whose corresponding private key is accessible only to a certain state and platform. Such data is inaccessible to any other state and platform. If the data is successfully decrypted, this indicates that the state the key was sealed to is the actual state of that machine.

Unlike the trusted TPM, the harddisk is not trusted. The harddisk merely provides bulk data storage, and no special properties are attached to it. Data may be modified on harddisk, and anyone can freely read the disk's contents. The Storage Service applies the necessary operations to provide us with persistent trusted storage.

## B. Attestation Service

The attestation service deals with providing metrics about the current configuration to remote parties.

We propose an way of proving a trusted state  $\sigma$  that does not directly involve invocation of the TPM. This has the advantage of not requiring the mere lookup of a configuration to already make use of the TPM. Instead, each participant in the Grid creates an *attestation token*  $\tau$ . Such a token is a tuple  $\tau := (p_{AIK}, p_K, \text{cert}(AIK), \text{cert}_{AIK}(p_K))$ .

This tuple consists of the following elements:  $\text{cert}_{AIK}(p_K)$  is a certificate (`KeyInfo` structure, to be exact) of the key including the metrics `config` of the reported state  $\sigma'$  and usage restrictions of

the key as detailed in Section III-A. `cert` (AIK) is the certificate of the AIK used. From `config`, `U` is able to distinguish a trusted  $\sigma'$  from an untrusted one, since the values uniquely identify a set of programs that have been loaded since booting the platform, and possibly also the state of certain critical configuration files. The simplest way of distinguishing trusted values from untrusted ones would be to keep a list of known-good values, and not consider any other values to be trusted. `certAIK( $p_K$ )` assures us that  $s_K$  can be used only in the reported state  $\sigma'$ , since the certificate identifies the key  $K$  as being sealed to `config`. `U` can make its trust decision “offline”, by examining the `config` contained in the  $\tau$ . If this `config` indicate a trusted  $\sigma'$ ,  $s_K$  will be accessible only if `P` still is in the same configuration. Since the token does not change over time, it can be distributed to other parties. If the  $\sigma$  of `P` ever changed, the  $\tau$  would automatically become invalid, although an explicit revocation might still be beneficial. Further details of this attestation mechanism and its security will be discussed in Section IV.

### C. Compartment Management Service

This component creates virtual machines (compartments), which run on top of the TCB. These virtual machines each comprise a portion of memory, and could (potentially) comprise a full-fledged operating system. Each of them is a full virtual machine, thereby enforcing the isolation properties necessary to prevent other compartments from accessing the internal information of a Grid job. That is, processes running in one compartments cannot access memory mapped to another. Thus, they are not able to interact with one another, except for well-defined interfaces. This isolation property guarantees both integrity and confidentiality. By setting up compartments in accordance to this isolation policy, this service actually provides the trusted virtualization we require. An approach to implementing these properties is provided in Section VI.

Furthermore, it keeps track of the identity of compartments, assigning to each compartment a unique identifier (UID). These UIDs are persistent for each compartment. The compartment management service is local to one machine, and does not attempt to track virtual machines across the Grid. Any communication between `P` and other participants needs to take place solely via interactions of the job itself. Being agnostic to the Grid and the particulars of its partition, the Compartment Management service is purely concerned with local secure execution.

### D. Storage Service

The Storage component provides trustworthy non-volatile storage based on an untrusted hard disk. Trusted Storage means that data stored by one compartment in one configuration is retrievable only by that compartment, in the same configuration, even if the machine has entered an untrusted state in the meantime. No other compartment or configuration is able to retrieve the information, or alter it undetected.

Confidentiality often implies not revealing the *existence* of certain data to unauthorized parties, not merely protecting its contents from access. For example, the mere existence of a plan to develop a product may be considered a business secret. Therefore, each compartment is outfitted with a separate namespace, which no other compartment may access (in addition to potentially shared storage). Any compartment is only allowed to access its own, and a potential shared namespace, thus preventing any other virtual machine from reading its confidential data.

While such separate namespaces protect data in the trusted state, `m` may enter an untrusted configuration, in which it can freely access all data. In order to prevent such an untrusted state from accessing or altering data, we utilize TPM sealing. By sealing a signature key, and using it to sign all stored data and the position on disk it was stored to, we can protect integrity (by preventing any but the trusted state from modifying data undetected). Alternatively, an Asymmetric key could be encrypted by a sealed key, and used to encrypt and MAC [7]. This would provide similar guarantees, with performance very similar to current whole-disk encryption applications. The same facilities (encapsulated in `load()` and `store()`) is made available to both TCB components, and virtual machines.

### E. Grid Management Service

This component handles the actual Grid job submission. It is responsible for receiving jobs, checking their access, and instantiating them. It will use the Compartment Management Service to create a private compartment for each job. If the job needs any special processing before becoming ready to be executed, this processing is done by this component, as well. Once such processing is done, a virtual machine image has been created from J, which can then be booted by the Compartment Management Service. The Grid Management Service furthermore takes the policy of the User, and notifies an enforcement component (not shown) of the restrictions and rights declared therein.

Its `submit()` interface also handles the freshness verification of the  $\tau$  available from the attestation component. For details regarding the protocol to receive jobs, see Section IV.

### IV. SCALABLE OFFLINE ATTESTATION

Attestation, in general, is the process of securely reporting the configuration (state) of a party to a remote challenger by cryptographic means. The most commonly discussed attestation variant requires a remote challenger to provide a random N, which is then signed (together with a hash over a subset of the current PCR values) by the TPM using an AIK. Since freshness is achieved by means of a random nonce, each interaction necessitates a new attestation (and thus, a new TPM-generated digital signature). However, generating signatures by TPM may significantly impact interaction speeds. Also, TPM commands can generally not be processed in parallel, introducing an undesirable bottleneck.

Additionally, without appropriate countermeasures, this technology could potentially be vulnerable to a race between a successful attestation, and a change of state before further interactions depending on the trusted state. If the state of the system changed after attestation concludes, yet before any further interactions take place, this change would not be noticed by the remote party. Also, without connecting attestation to a PKI identity, an attestation challenge could in principle be replayed by an attacker to a trusted platform, with the attacker forwarding the platforms reply to its own verifier.

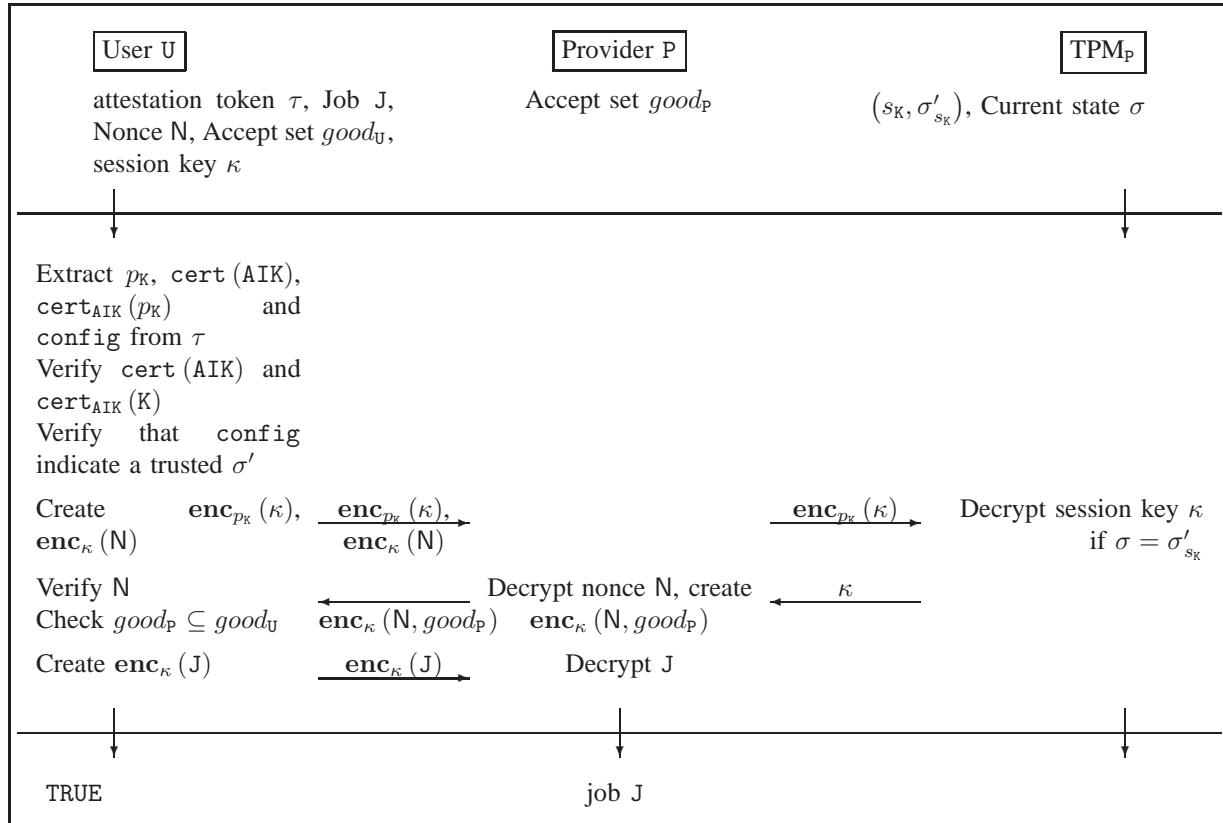


Fig. 3. Submission Protocol `submit()`

Scalable offline attestation, as previously mentioned, is intended to enhance some aspects of current attestation systems. Having an attestation token which can be freely distributed among the virtual organization as an informational item is advantageous, since this token states the current configuration of a provider P, without forcing the prospective user to interact with that provider right away. The user can collect such tokens over time, and select the most appropriate configuration offline.

Naturally, since such a token cannot provide any guarantee of freshness of the reported values, some verification needs to occur, otherwise the attestation would be meaningless. We propose to follow a *sealed key* approach, where a private key is only accessible to the TPM (and therefore, the provider) if the provider is in the same state the key was stored in. This naturally partitions verification of P’s state into two phases: Token creation and freshness verification.

A provider P creates an attestation by using its TPM. The attestation service instructs the TPM to create a non-migratable key sealed to a collection of PCRs. Then, the attestation service uses the TPM’s CertifyKey operation to create a certificate  $\text{cert}_{\text{AIK}}(p_K)$  with an AIK. This certificate is a KeyInfo structure signed by an AIK as described in Section III-A. The attestation service then constructs the attestation token  $\tau$  from the public key  $p_K$ , the certificate of this key  $\text{cert}_{\text{AIK}}(p_K)$ , the public part of the AIK  $p_{\text{AIK}}$ , and a certificate of the AIK  $\text{cert}(\text{AIK})$ .  $\tau$  is then published by the attestation service.

Freshness verification is addressed by the protocol in Figure 3, which includes the actual submission of the job. The user generates a symmetric “session key”  $\kappa$ , which is encrypted using  $p_K$ . The session key can be decrypted by the TPM of P, if its state still equals the state when creating the  $\tau$ , that is, its reported state  $\sigma'$ . The rationale for including this key is twofold: First, asymmetric cryptography is slower than symmetric methods by orders of magnitude, making the establishment of a “load key” a common technique. Second, this reduces the necessary TPM operations to a single asymmetric decryption, which is *at worst* not slower than the signature utilized by standard techniques, since the same underlying cryptographic operation is used.

The submission protocol further guarantees *transitive* trust. That is, as the job or parts of the job get delegated from a provider to sub-providers, it is assured that each party that will be entrusted with the job’s data will satisfy the original submitter’s requirements. This is done by ensuring that each platform  $X$  that gains control of user U’s job J must satisfy  $\text{good}_X \subseteq \text{good}_U$ .

Publishing an attestation token  $\tau$  in effect becomes an advertisement stating that a certain state  $\sigma$  will be maintained. If the  $\sigma'$  contained in the token is considered good by the user U, that user can utilize the key  $p_K$  to encrypt data. This data is unrecoverable for any node other than U, and furthermore, can only be accessed if  $\sigma_P = \sigma'_{s_K}$ . Thus, verification of P’s ability to access  $s_K$  is sufficient to ensure that  $\sigma_P = \sigma'_{s_K}$ . The assurance of  $s_K$  actually being accessible only in state  $\sigma'$  is crucial to this protocol. If it were not guaranteed,  $s_K$  might be accessible in some malicious state  $\sigma$ , and all the guarantees derived from the trustworthiness of the TCB would no longer be valid. The certificate generated by the TPM assures this property of the key, by signing it with a TPM-internal AIK. This key cannot be misused by the platform owner, and the stated PCR values of this certificate are correct.

## V. SECURITY ANALYSIS

### A. Security of Offline Attestation

The offline attestation mechanism proposed in Section IV is secure against man-in-the-middle attacks: If the user U seals a job to a trustworthy attestation token  $\tau$ , only the platform in possession of the private part of the key  $p_K$  from  $\tau$  can unseal the job, and only if it is in the state indicated by  $\tau$ . An adversary cannot decrypt the job, even if it is running on the platform with the TPM that holds the private key, if its state  $\sigma$  does not equal the  $\sigma'$  contained in  $\tau$ . Conventional techniques need to include additional verifications (such as tying an AIK to a PKI identity) to achieve the same assurance, thus demonstrating an advantage of our proposal.

Delegation with transitive trust ensures that every provider P that gets a job J can only access J if it is in a state  $\sigma$  that is trusted by the original submitter U, i.e.,  $\sigma \in \text{good}_U$ . Transitive trust is achieved during delegation without communication with the submitter, because the provider that wishes to

transfer a job offline-attests other providers before transmitting the job. The delegating provider  $P_1$  acts as user of the new provider  $P_2$  and verifies that  $good_{P_2} \subseteq good_{P_1}$ , which immediately implies that  $good_{P_2} \subseteq good_U$ . Hence, the policy of the new provider  $P_2$  is also acceptable to the original user.

Moreover, offline attestation is secure against replay attacks, under the assumption that state changes can only occur between protocol runs. Replaying of old, trustworthy attestation tokens does not help an adversary: the TPM will not allow decryption if the current PCR values do not match the values the key was sealed against.

However, like conventional attestation, our protocol is vulnerable to TPM compromises. A compromised TPM can expose the secret key to an adversary, which enables that adversary to attest to arbitrary states. Revocation of AIKs must be possible to limit the damage caused by this attack. Like conventional attestation, another risk of offline attestation is corruption of the running TCB. If an adversary can corrupt the TCB while the system is running, it could change the system's state  $\sigma$  without changing the PCRs (and thus, deviating from  $\sigma'$ ).

### B. Integrity Protection

Since we can establish a secure (confidential and integrity-protected) channel from user  $U$  to provider  $P$  using standard tools such as TLS, we need not consider in-transit modifications. Thus,  $P$  receives, for the purpose of this analysis, an unaltered job  $J$ . We need to consider two kinds of integrity requirements for that image: before being instantiated, and while executing. Since results are reported directly, their integrity can once more be achieved by established solutions.

If job execution is delayed by the Grid Manager, the job image and policy are stored in trusted storage. The Storage Management Service's key is stored sealed, which guarantees that access to it is granted only to the same job in the same system state. In an untrusted state, no access is granted. Therefore, if a piece of data  $X$  is altered, the signature of that item cannot be updated, and the modification is detected during the next load () operation.

While the job  $J$  is executing, the isolation properties of our system guarantee that no untrusted application can gain access to the memory regions assigned to  $J$ . Since access is forbidden, integrity is guaranteed. Circumventing such barriers would necessitate breaching the TCB. As the TCB is based on a virtualization layer, even attack scenarios like "blue pill"<sup>1</sup> are ineffective, since such rootkits can only virtualize conventional systems that do not use virtualization techniques themselves. However, even if such a system were able to virtualize a virtualization layer, it would either need to be loaded before the TGA (and thus, be measured in the trusted boot process), or compromise the TCB.

### C. Confidentiality Protection

Similar to integrity, the two mechanisms employed for stored data and in memory data also protect confidentiality. The Compartment Manager enforces isolation between the VMs it creates, foiling in-memory eavesdropping, i.e., one process accessing data inside the virtual memory of another process. Sealing prevents untrusted configurations from decrypting data stored in non-volatile storage.

Breaching confidentiality implies breaching the TCB for the in-memory scenario, since the TCB enforces virtualization and therefore, limits each application to its own VM, while decrypting stored data outside of a trusted state would necessitate breaking the encryption scheme used, which we likewise consider infeasible.

## VI. DISCUSSION AND RELATED WORK

### A. Integration of Legacy Systems

Grid toolkits, such as Globus [11], [13], offer platform abstraction, authentication and other utility services. While designing our proposal, it was our intention to maintain interoperability with these solutions whenever possible. Towards this end, we aim to provide a way to continue using Grid applications designed for these environments, without giving up the advantages our architecture offers.

<sup>1</sup>Presented by Joanna Rutkowska at Syscan '06, <http://theinvisiblethings.blogspot.com/>

One possible way for such an integration would be to provide, for each supported toolkit, an executable image featuring a “default installation”. Whenever an invocation for a service using that toolkit is received, it is instantiated, and the request forwarded to that instance.

However, the Grid toolkit must be part of the TCB. After all, a malicious provider might use a good base configuration, and put all its attack code into a modified toolkit image. The Attestation Token  $\tau$  should contain measurements of all execution environments available as “default installations” on the platform. Thus, the benefits of our proposal become applicable, without forcing the user to significantly change its use of the Grid.

Alternatively, a Grid Job may consist of a full, bootable image of a filesystem. While this is a radically different approach from traditional Grid methods, this technique demonstrates the added isolation properties much more prominently. Additionally, this approach does not imply further trusted code, which is desirable in light of keeping the TCB small, and of low complexity.

If neither option is appropriate, and interaction with an untrusted system becomes necessary, of course Trust Asymmetry becomes an issue once again. While an untrusted provider can not be relied on to handle confidential data in a sensible manner, utilizing a system similar to the proposal in [18] allows us to “filter” untrusted input data, and retain some integrity assurances for the entire system, even if it partially depends on untrusted (low-integrity) inputs.

### B. Possible Realization

Our approach to implementing the abstract components of the TGA architecture is based on the PERSEUS framework [24], [26], [28]. Another approaches similar to PERSEUS is Nizza [16]. Due to space restrictions, we will not describe it in detail, but instead will provide these details in a future contribution.

Several demonstrators have already been implemented within the PERSEUS framework, such a secure VPN service, and harddisk encryption [29], [4], as well as a secure home banking demonstrator, all as part of the ongoing research and development projects Open Trusted Computing (OpenTC, <http://www.opentc.net>) and European Multilateral Secure Computing Base (EMSCB, <http://www.emscb.org>, see also [28]).

PERSEUS’ current implementation is based on a micro-kernel with paravirtualized Linux. Its interfaces are generic, however, and should not pose significant problems to port to other architectures (such as Xen). Its design allows for a strong separation of responsibilities even among the TCB (by running services as separate compartments), and makes verification of software significantly simpler than a monolithic approach.

### C. Attestation

In this paper, we have assumed binary (hash-based) attestation. The *good* set of a host is simply modeled as the set of configuration values (as encoded by PCR registers) that represent “good” machine states. While this is a very simple and straightforward mapping, and binary attestation techniques can form a basis for verifying and authenticating platforms ([31], [15]), current techniques suffer from some shortcomings which make these systems hard to use in practice, unless further techniques are employed. One such limitation to binary attestation schemes is their fragility. Minute changes in the verified files, which may not affect their semantics in any way, will lead to an invalidation of the stored hash, and thereby to the rejection of the state.

Our proposed protocol does not feature any privacy guarantees, unlike DAA [9]. In fact, the sole purpose of our  $\tau$  is to distribute configuration information ahead of time. Other works consider privacy to be a major issue in attestation, and integrating privacy with our work might be a worthy avenue of future research.

Usually, the verifier intends to check *properties*, rather than binary code details. Therefore, it is sensible to also base attestation on properties. [27], [25] offer such a property-based attestation scheme, which additionally supports a higher level of privacy for the attestee than binary methods would support. Especially in heterogeneous Grid environments with a plethora of acceptable configurations, this approach should be advantageous.

#### D. Alternative and Complementary Approaches

Several authors have suggested methods to increase reliability of Grid computation without utilizing TC technology. Techniques such as task replication, or the suggestion to introduce Quiz tasks [38] to detect misbehaving providers have aimed to protect result integrity. However, these techniques are wasteful in terms of resources, and often not very resistant to the presence of multiple colluding adversaries.

We have considered security on the level of the virtualization layer, without going into details about realizing trust and measurements within virtual machines. Sailer et al. [30], [22] have conducted an inquiry into the possible enforcement of MAC policies at this level, with encouraging results. In [32], they also conducted some inquiries into a possible measurement architecture. While these results, only concern themselves with the measurements of one virtual machine, as opposed to the entire system, they still may be useful in measurement and reporting of VM states.

The work of Mao, Jin and Martin [20] on integrating Grid technologies with Trusted Computing already contained several valuable contributions on both increasing security and reducing overhead by utilizing TC technology. These advances can be integrated into our proposal, resulting in an overall more secure grid, which can achieve comparable or even increased performance to current practice.

Encrypted computation [3], [33] could possibly replace trusted computing for some problems. By processing encrypted data without decrypting it, some tasks can be completed, without ever revealing plain text. Unfortunately, however, these techniques have limited use outside the domain of some algebraic problems, and widespread adaption of these techniques seems unlikely.

### VII. CONCLUSIONS AND FUTURE WORK

In this paper, we have laid the groundwork for a Grid architecture that is better-suited to execute sensitive jobs on untrusted hosts than current practices. By enforcing isolation properties, such a job can utilize the computation power of a host, without having to risk exposing data to the platform owner. Our approach is likely to require some (efficient) re-engineering of current grid frameworks, but allows for techniques very similar to current practices. In effect, a layer of negotiation can be added “on top” of existing solutions, which can still be used to operate the working load of the Grid, once trust has been established.

Interoperability between Grid Manager services needs to be addressed. If every trusted OS offered a different interface, interoperability would needlessly suffer. Therefore, standards on the particulars of job submission and attestation token retrieval need to be agreed upon.

Scalable offline attestation still necessitates direct interaction with the TPM in the submit phase, thus limiting its uses. We are currently investigating the removal of this dependency, possibly by storing the secret key not in the TPM, but under the control of the Storage Management Service. The Storage Management Service can, if strong isolation is achieved, fulfill the same functions as the TPM, and allow access to the key only in the state the key was created in. However, due to the more complex trust model, and greater avenues for attack, we are still investigating this option.

Further work is needed in realizing the trusted grid architecture. Some of the components we identified have not yet been fully examined. Also, issues of reusing and integrating existing infrastructure is work in progress together with members of the Grid community.

### REFERENCES

- [1] AMD Pacifica Technology Home Page. [http://developer.amd.com/assets/WinHEC2005\\_Pacifica\\_Virtualization.pdf](http://developer.amd.com/assets/WinHEC2005_Pacifica_Virtualization.pdf).
- [2] Intel LaGrande Technology Home Page. <http://www.intel.com/technology/security/>.
- [3] J. Algesheimer, C. Cachin, J. Camenisch, and G. Karjoth. Cryptographic security for mobile code.
- [4] A. Alkassar, M. Scheibel, A.-R. Sadeghi, C. Stüeble, and M. Winandy. Security architecture for device encryption and vpn. *Accepted for ISSE 2006*.
- [5] F. Azzedin and M. Maheswaran. Towards trust-aware resource management in grid computing systems. In *CCGRID'02*, 2002.
- [6] J. Basney, W. Nejdil, D. Olmedilla, V. Welch, and M. Winslett. Negotiating trust on the grid. In *2nd Workshop on Semantics in P2P and Grid Computing*, New York, May 2004.

- [7] M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. *Lecture Notes in Computer Science*, 1109, 1996.
- [8] R. Boutaba and A. Polyakis. Towards extensible policy enforcement points. In *LNCS*, volume 1995, page 247ff, 2001.
- [9] E. Brickell, J. Camenisch, and L. Chen. Direct anonymous attestation. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 132–145, 2004.
- [10] P. A. Dinda. Addressing the trust asymmetry problem in grid computing with encrypted computation. In *LCR '04: Proceedings of the 7th workshop on Workshop on languages, compilers, and run-time support for scalable systems*, pages 1–7, New York, NY, USA, 2004. ACM Press.
- [11] I. Foster. Globus toolkit version 4: Software for service-oriented systems. In *IFIP International Conference on Network and Parallel Computing*, 2005.
- [12] I. Foster, C. Kesselman, G. Tsudik, and S. Tuecke. A security architecture for computational grids. In *Fifth ACM Conference on Computers and Communications Security*, Nov 1998.
- [13] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the grid: Enabling scalable virtual organizations. *International Journal of Supercomputer Applications*, 3(15), 2001.
- [14] A. Fuggetta, G. P. Picco, and G. Vigna. Understanding code mobility. *IEEE Transactions on Software Engineering*, 24(5), May 1998.
- [15] T. Garfinkel, B. Pfaff, J. Chow, M. Rosenblum, and D. Boneh. Terra: A virtual machine-based platform for trusted computing. In *In Proceedings of the 19th ACM Symposium on Operating Systems Principles*, 2003.
- [16] H. Hertig, M. Hohmuth, N. Feske, C. Helmuth, A. Lackorzynski, F. Mehnert, and M. Peter. The nizza secure-system architecture. In *CollaborateCom*, 2005.
- [17] K. Hwang, Y.-K. Kwok, S. Song, M. C. Y. Chen, Y. Chen, R. Zhou, and X. Lou. Gridsec: Trusted grid computing with security bindings and self-defense against network worms and ddos attacks. In *LNCS 3516*, pages 187–195, 2005.
- [18] T. Jaeger, R. Sailer, and U. Shankar. PRIMA: policy-reduced integrity measurement architecture. In *SACMAT '06: Proceedings of the eleventh ACM symposium on Access control models and technologies*, pages 19–28, New York, NY, USA, 2006. ACM Press.
- [19] U. Kühn, K. Kursawe, S. Lucks, A.-R. Sadeghi, and C. Stübke. Secure data management in trusted computing. In *Proceedings of the Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, number 3659 in LNCS, pages 324–348, 2005.
- [20] W. Mao, H. Jin, and A. Martin. Innovations for grid security from trusted computing. Made available online at [http://www.hpl.hp.com/personal/Wenbo\\_Mao/research/tcgridsec.pdf](http://www.hpl.hp.com/personal/Wenbo_Mao/research/tcgridsec.pdf).
- [21] P. Mazzoleni, B. Crispo, S. Sivasubramanian, and E. Bertino. Efficient integration of fine-grained access control in large-scale grid services. *Proceedings of the 2005 IEEE Conference on Secure Computing Services*, 01:77–86, 2005.
- [22] J. M. McCune, S. Berger, R. Cañeres, T. Jaeger, and R. Sailer. DeuTeRiuM – a system for distributed mandatory access control. Available online at <http://domino.research.ibm.com/library/cyberdig.nsf/1e4115aea78b6e7c852%56b360066f0d4/eac6708db1c54cf38525710f005b7aff?OpenDocument>, accepted for presentation at *ACSAC 2006*.
- [23] R. Nabhen, E. Jamhour, and C. Maziero. A policy based framework for access control. In *ICICS*, pages 47–59, 2003.
- [24] B. Pfitzmann, J. Riordan, C. Stübke, M. Waidner, and A. Weber. The PERSEUS system architecture. Technical Report RZ 3335 (#93381), IBM Research Division, Zurich Laboratory, Apr. 2001.
- [25] J. Poritz, M. Schunter, E. van Herreweghen, and M. Waidner. Property attestation: Scalable and privacy-friendly security assessment of peer computers. Early dissemination research report.
- [26] A.-R. Sadeghi and C. Stübke. Bridging the gap between TCPA/Palladium and personal security. Technical report, Saarland University, Germany, 2003.
- [27] A.-R. Sadeghi and C. Stübke. Property-based attestation for computing platforms: caring about properties, not mechanisms. In *NSPW*, pages 67–77, 2004.
- [28] A.-R. Sadeghi, C. Stübke, and N. Pohlmann. European multilateral secure computing base - open trusted computing for you and me. *Datenschutz und Datensicherheit DuD, Verlag Friedrich Vieweg & Sohn, Wiesbaden*, 28(9):548–554, 2004.
- [29] A.-R. Sadeghi, C. Stübke, M. Winandy, and M. Scheibel. Design and implementation of a secure linux device encryption architecture. Technical report, The EMSCB Project, 2006. available at <http://www.prosec.rub.de/hddenc.pdf>, accessed 10.06.2006.
- [30] R. Sailer, T. Jaeger, E. Valdez, R. Caceres, R. Perez, S. Berger, J. L. Griffin, and L. van Doorn. Building a mac-based security architecture for the xen open-source hypervisor. In *ACSAC '05: Proceedings of the 21st Annual Computer Security Applications Conference*, pages 276–285, Washington, DC, USA, 2005. IEEE Computer Society.
- [31] R. Sailer, T. Jaeger, X. Zhang, and L. van Doorn. Attestation-based policy enforcement for remote access. In *CCS '04: Proceedings of the 11th ACM conference on Computer and communications security*, pages 308–317, New York, NY, USA, 2004. ACM Press.
- [32] R. Sailer, X. Zhang, T. Jaeger, and L. van Doorn. Design and implementation of a tcb-based integrity measurement architecture. IBM Research report, published online at <http://byte.csc.lsu.edu/~durresi/7502/reading/rc23064.pdf>. To appear at the *13th annual USENIX Symposium*.
- [33] T. Sander and C. Tschudin. Protecting mobile agents against malicious hosts. In *Lecture Notes in Computer Science*, volume 1419, pages 44–60, 1998.
- [34] M. Smith, T. Friese, M. Engel, and B. Freisleben. Countering security threats in service-oriented on-demand grid computing using sandboxing and trusted computing techniques. page (to appear), 2006.

- [35] S. Song, K. Hwang, and Y.-K. Kwok. Trusted grid computing with security bindings and trust integration. *Journal of Grid Computing*, (3):53–73, 2005.
- [36] S. Song, Y.-K. Kwok, and K. Hwang. Security-driven heuristics and a fast genetic algorithm for trusted grid job scheduling. In *19th IEEE International Parallel and Distributed Processing Symposium*, 2005.
- [37] M. R. Thompson, A. Essiari, K. Keahey, V. Welch, S. Lang, and B. Liu. Fine-grained authorization for job and resource management using akenti and the globus toolkit. *Computing in High Energy and Nuclear Physics*, 24–28, Mar. 2003.
- [38] S. Zhao, V. Lo, and C. Gauthier-Dickey. Result verification and trust-based scheduling in peer-to-peer grids. In *Proceedings of the Fifth IEEE International Conference on P2P (P2P2005)*, 2005.