Adaptive defense measures against the security hazards induced by systems virtualisation

Vitalian A. Danciu[1], Nils gentschen Felde[1], Wolfgang Hommel[1], Tobias Lindinger[1]

Munich Network Management Team, Munich, Germany
{danciu,felde,hommel,lindinge}@mnm-team.org

1 Motivation

The explosive proliferation of systems virtualisation supports a more adaptive delivery of IT services. In the area of security management, however, virtualisation carries beneficial as well as detrimental implications. Fig. 1 shows the sources of requirements on security management: they originate either from technical sources, in a bottom-up view, or from high-level goals. As the business management is held liable for financial damage caused by IT security incidents, e. g. due to the BASEL II and SOX regulations [1,2,3], security measures are no longer a task which is exclusively handled by the IT staff. Governance, Risk, and Compliance (GRC) frameworks focus the support of these three business critical management areas on the strategic level.

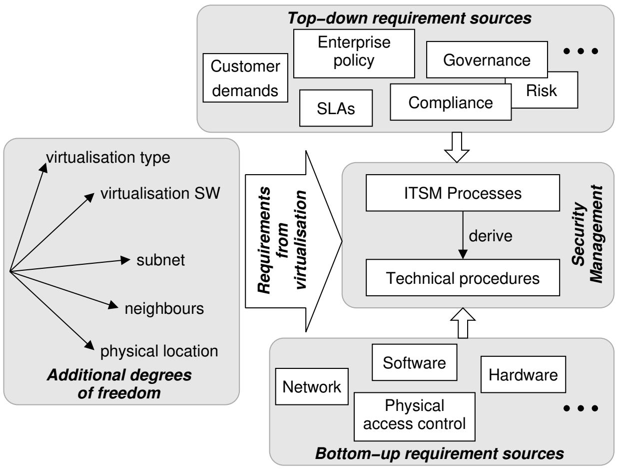

Fig. 1. Additional requirements

In this paper, we analyse and classify new hazards that are due to the additional degrees of freedom (shown to the left in Fig. 1) introduced by systems virtualisation. We consider technical, as well as process-related aspects and derive action items for handling the issues that originate from systems virtualisation.

Until recently, server machines could to some degree be viewed as separate entities providing services under a clear set of security constraints. Taking into account *VM migration*, i. e. the displacement of virtual machines between physical hosts, a number of assumptions common in traditional (i. e. non-system-virtualised) service provisioning no longer hold true: VMs may reside together with others (*neighbours*) on the same physical machine, change their *physical location* and attach to different *subnets*. While such changes are currently being actuated by (human) systems managers, the technological prerequisites are being made available to perform such VM re-deployment in an automated manner, e. g. in order to achieve availability or load balancing objectives [4]. Existing virtualisation strategies (full virtualisation, para-virtualisation, etc.) realise different *types* of capability sets, while critical security issues are being introduced by the virtualization software itself.

All things considered, the threats against services provided on (even partially) virtualised infrastructures call for amendments to technical, as well as organisational security practices. In the following section, we examine the change in security settings, taking into account two different perspectives (customer, process-oriented). According to the deliberations in Sec. 2, we propose protective measures in Sec. 3.

2 Threat assessment

A service-oriented view on security management in virtualised environments encompasses technical and organisational aspects, as well as the service customers' perspective. The additional degrees of freedom as well as top-down and bottom-up aspects lead to requirements influencing security management processes and the thereof derived technical procedures (see Fig. 1).

2.1 Avenues of attack

In the following, we examine risks introduced by the additional degrees of freedom in virtualised systems, taking into account the migration functions of VMs. Their influence is as follows: 1) *Physical location:* The control of physical access to computing hardware, power supplies, cooling facilities and so on is critical for IT security. Data centres apply access control measures at varying levels of security. Hence, if only service quality or load balancing aspects are considered, that protection is equal to the lowest level of protection of any data centre operating a VM. 2) *Neighbourhood:* A changing number and "quality" of neighbours may reside on the same system at a certain time. Meanwhile, as a function of the customer and type of service, customers may demand their VMs to be isolated from other customer's VMs, as we explain in more detail in Sec. 2.2. 3) *Subnet:* Services may rely on certain protection of the network or, conversely, some services may rely on relaxed network security settings for their own proper operation (e. g. certain open TCP ports). If the location of a service in the network is to be completely opaque, we would need to accept that this protection is granted to a varying degree, as given by the subnet. To enforce a certain security standard for a service, the virtualisation infrastructure needs to be aware of the position of the service in the network and restrict the service to proper locations. 4) *Virtualisation type:* Virtualisation products are designed according to different paradigms which typically constitute a trade-off between VM performance and totality of the abstraction mechanism. Depending on the approach employed in a product, a certain degree of direct access to physical hardware is granted to VMs. Thus, the level of isolation between VMs (and, hence the level of security) varies with the design or *type* of virtualisation platform. 5) *Virtualisation software:* The brand of virtualisation software, its version as well as different patch-levels of the systems entail different levels of security, as is the case with any systems software. Already, efforts are under way to enable migration of VMs between different virtualisation platform *types*, supplied by different software vendors.

2.2 Customer view

The individual mode of hosting a service, as well as the subservices upon which it depends, influence its *threat environment*. Hence, to ensure an adequate level of security and data protection, a customer will be interested in contractually adapting the following service level parameters: 1) *number of neighbours* co-located at a given time; 2) *changes to the co-location*, i. e. "new neighbours"; 3) *identity/profile of neighbours*, in order to determine whether their presence constitutes a potential risk. For example, business critical production systems of one customer typically must not be hosted along with test machines of another customer on the same server. 4) *level of security maintained by neighbours*, in order to determine how easily co-located services or resources could be compromised; 5) *track record of neighbours*, i. e. successful attacks on their services, or attacks originating on their resources.

It is foreseeable that customers, which have demands on security due to the nature of the service or to fulfil compliance regulations (compare Sec. 1), will ask information with respect to these parameters and demand that its provisioning is formalised in the service level agreements. Additional monitoring, logging and reporting will be necessary in order to satisfy such demands.

2.3 ITSM view: Information base and processes

Research and practice have proven that ITSM processes utterly require suitable tool support, and that all processes depend on accurate management information. To take into account virtualisation aspects, information schemas need to differentiate between physical and virtual machines. In addition, they must respect the opaque location of service resources (a consequence of introducing VM migration) in dependency graphs between service and service resources.

Process definitions, often derived by means of reference processes originating in best practices collections, need to be adapted in order to reflect the added requirements discussed above: 1) *activity extensions* to reflect the additional monitoring and reporting aspects; 2) *standard paths* for the recovery of compromised services provisioned by means of virtual systems, including new interfaces to IT forensics; 3) *service level and service catalogue amendments* to include specifically provisioning based on VM technology.

A brief overview of such extensions is discussed in the following section.

3 Virtualisation security related adaptation of ITSM processes

Virtualisation per se has an obvious impact on many ITSM processes, such as capacity management. Similarly, the security management process affects all the other ITSM processes. This section discusses the intersection of these aspects, i. e. which of the service delivery and service support processes are influenced by the security properties, or lack thereof, of service virtualisation.

We first investigate *Configuration management* (CM), as it is the linchpin when discussing the adaptations of all the other ITSM processes. CM is tool supported by a database (CMDB) which stores information about so-called configuration items (CIs) and their interrelationships. This management information is utterly required by other processes, e. g. in order to determine the security impact of incidents or planned modifications. Virtualisation implicates the following security relevant extensions:

- With both hosts and VMs being CIs, we propose a new type of relationship "*VM v_i is currently running on host h_j*", which is required to quickly assess the impact of an security incident at v_i on the other VMs running on h_j.
- We propose that the CMDB must be updated in real-time to properly reflect the current physical location of VMs. This way, the reporting and other service level management tools, which already have interfaces to the CMDB, can be used to regularly check the compliance of the current VM distribution with SLAs and internal security policies.
- CMDB baselines, i. e. enterprise-internal standards for sets of configurations, must be adapted to the VM templates used for rapid VM deployment and updating. Ensuring that e. g. the latest security patches are included in these baselines will become one of the primary security management goals.
- New virtualisation specific and security related policies will be subject to CM: For example, we propose that *migration policies* define triggers and conditions for DRS (dynamic resource scheduling) by means of rulesets which restrict to which host a VM may be moved.

Incident management must differentiate between physical machines and VMs for any security event, such as those reported by intrusion detection systems. An affected VM may need to be isolated depending on the severity level of the incident, e. g. by triggering its migration to a dedicated quarantine server. Automated snapshots, which cannot be manipulated from within the VM, can be created to enhance later IT forensics even if the attacker attempts to remove her traces.

Problem management, when triggered by security incident management, needs to assess the threats to the VM's present and previous neighbourhoods since the incident was reported. For example, if an attacker has compromised a VM, not only the same vulnerability might affect identically configured other VMs, but exploits targeting the virtual machine monitor (VMM) might have lead to the compromise of neighbouring VMs as well.

Change management must honour the virtualisation specific technical security countermeasures; for example, migrating a VM might require an appropriate update of the target server's anti-spoofing configuration. The available protection measures lead to the necessity of new types of pre-authorised changes in order to apply them automatically. This in turn requires that virtualisation technology specialists must become members of the enterprise's Change Advisory Board (CAB).

Release management faces the challenge of arbitrarily changing VM neighbourhoods. While snapshots greatly simplify rollback planning, release management must ensure the compatibility of authorised changes with the actual VM instances, which requires that testing new releases considers the migration policies discussed above.

Availability management and *capacity management* must take into account that security measures may necessitate a temporary migration of VMs to other hosts. Enough spare resources must be provided to cope with potentially several parallel attacks against the virtualised infrastructure. As an attack against a host affects multiple VMs and services, leading to the instant violation of multiple SLAs, traditional formulas and tools for spare resource planning are no longer sufficient. Trends in the resource demand based on the enterprise's specific threat profile can be derived from the history stored for each CI in the CMDB.

Security management itself has many new options at its disposal like e. g. easy to create disk images and RAM dumps. On the other hand, security management obviously becomes more complex because new specific attacks against virtualised infrastructures are possible, while none of the traditional attacks cease to exist. New virtualisation specific security policies need to be defined, and service desk personell as well as system and service administrators need to be trained and sensitised to the VM specific security properties. Additionally to the security measures deployed on each VM, which mostly are the same as if it were a physical server, security measures and monitoring must be set up for the hosts.

Already this short summary of virtualisation security implications makes it obvious that the adoption of virtualisation technology has a major impact on ITSM processes that has not been accounted for in current best practice collections.

4 Observations in a real-world environment

The MNM-Team teaches a practical course on IT security for graduate students of two of the Munich universities, LMU and TUM. Recently, the complete lab environment has been migrated to a Xen-based virtualisation solution. In the curriculum, various security flaws are explained, and their misuse is illustrated in experiments using sniffers, port scanners, several "hacking" tools, and executing Denial-of-Service (DoS) attacks between VMs hosted by the same Virtual Machine Monitor.

This lab environment is a perfect sandbox to illustrate the virtualisation security aspects discussed in this paper: While students benefit from the virtualised infrastructure because they can work on it remotely at home, we can closely monitor attacks against the infrastructure originating from the internet. Furthermore, in each semester, students try to use the latest exploits and DoS tools to hack each other, which gives detailed insight, e. g. into resulting performance issues, and grants us the opportunity to try out VM-vs-VM intrusion detection mechanisms in a realistic environment.

5 Prospect

Virtualisation security is a new research area motivated by the new types of attacks introduced along with the new technology. This paper sketches an overview of our structured approach towards dealing with these emerging challenges. On the one hand, it focuses on the classification of attack types and their technical counter-measures, backed by the outlined scenario. On the other hand, we presented several concrete amendments to IT management reference processes as described in the well-known ITILv2 best practices collection.

Acknowledgment

The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful discussions and valuable comments on previous versions of this paper. The MNM-Team directed by Prof. Dr. Heinz-Gerd Hegering is a group of researchers of the University of Munich, the Munich University of Technology, and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. Its web-server is located at http://www.mnm-team.org.

References

1. Bank for International Settlements: Basel II: International Convergence of Capital Measurement and Capital Standards: a Revised Framework (2004)
2. Sarbanes, P., Oxley, M.: H.R.3763 Sarbanes-Oxley Act of 2002 (2002)
3. Bundesamt für Sicherheit in der Informationstechnik (BSI): BSI-Standard 100-2: IT-Grundschutz Methodology. Technical report, Bundesamt für Sicherheit in der Informationstechnik (BSI) (2005)
4. Clark, C., Fraser, K., Hand, S., Hansen, J.G., Jul, E., Limpach, C., Pratt, I., Warfield, A.: Live migration of virtual machines. In: Proceedings of the 2nd ACM/USENIX Symposium on Networked Systems Design and Implementation (NSDI). (2005) 273–286