Host virtualization: a taxonomy of management challenges

Vitalian A. Danciu

Munich Network Management Team
Ludwig-Maximilians-Universität, Munich, Germany
danciu@mnm-team.org

Abstract: Host virtualization is quickly being introduced to production environments as it facilitates the recent years' computing centre consolidation efforts. While its introduction offers new opportunities in IT management, it also presents challenges that are yet to be tackled. In this paper, we chart these areas of concern according to established conceptual management frameworks and juxtapose the result to a survey of current work.

1 Introduction

The proliferation of production-grade host virtualization solutions in recent time has been welcomed by data centre operators for the promises offered with respect to cost savings and management facilitation. Applications, services, as well as Operations and Business Support Systems software (OSS/BSS) are being increasingly provided from a virtual machine platforms. Thus, the management of virtual hosts is a prerequisite for the operation of telecommunications networks and services.

It appears, however, that the (significant) benefits derived in the short term from the introduction of virtualization obscure the management requirements that arise from its operation. Figure 1 shows the reliance of the provided services, as well as support and management systems on each other, and on the increasingly virtualized IT resources. This paper discusses the additional challenges to technical and process-oriented service management incurred from the operation of host virtualization technology.

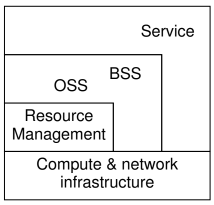

Figure 1: Dependence on infrastructure

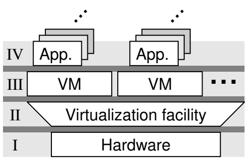

Figure 2: Fundamental architectural layers in host virtualization

Suppliers of virtualization facilities continually extend their products with new functions. Different strategies (OS virtualization, emulation, para-virtualization etc, see e.g. [NC05, SN05, Wol07]) are being employed for realising the *basic virtualization function*, striking a balance on benefit tradeoffs regarding transparency, performance and portability/ease of implementation. In principle, however, any host virtualization approach will adhere to the architectural layers sketched in Figure 2. A virtualization facility (II) (e.g. a modified operating system kernel, an emulator, a virtual machine monitor) is executed on a physical machine (I), in order to deploy and run multiple instances of one or more operating systems (III). Thus, an additional abstraction layer is introduced between machine hardware and the application software (IV) providing service to customers/users.

While management of virtualization facilities and VMs has become a selling point for commercial offerings, existing management tools emphasize specific issues of virtualization management instead of pursueing an integrated approach; needless to say, vendor specific management is dominant. Conversely, virtual hosts are supported only in part by current management concepts (MIBs, mgmt functions, management processes). Integrated management systems do not yet handle VMs themselves, they rely on proprietary resource management systems. The long-term aim must be to embed the management of virtual hosts into the existing architectures and systems for integrated management.

Host virtualization forfeits certain assumptions that management concepts (and systems) tacitly rely upon [GKM$^+$03] and introduces new dynamic aspects to infrastructure and service management. To understand the management challenges originating from host virtualization, we must assess the differences introduced by this technology in different areas, as suggested in Figure 3: virtual host environments introduce new features that in turn allow for novel means of provisioning infrastructures for delivering service. These management scenarios for virtualization pose requirements on management concepts and systems. Vendors and products for host virtualization differ in the technology used, the strategy for offering the basic virtualization functions, in feature sets and management interfaces.

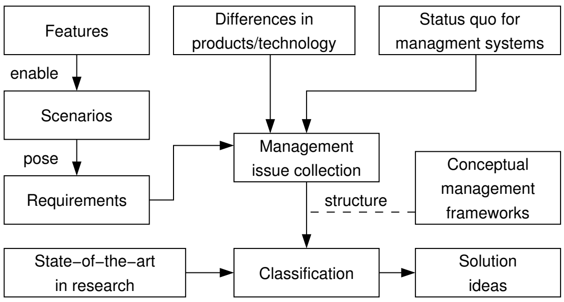

Figure 3: Methodology overview

Finally, management systems—resource management systems for VMs, as well as integrated management systems—differ in the possibilities they offer for managing environments containing virtual machines. These constitute sources for management issues.

The core contribution of this paper is to identify, analyse and classify the management issues introduced by host virtualization, and to offer ideas on solution alternatives. We employ established conceptual frameworks for IT management to order and classify these issues, and we discuss current research work that can be aligned to them. A number of solution ideas result for selected areas. We analyse management challenges organised along the characteristic features of host virtualization, identified in Section 2, and functional areas of management in Section 3 to assess the impact of the deployment of virtualization technology. The management facilities addressing each of these challenges require appropriate extensions or additions in order to be able to handle virtual hosts in the same manner as traditional, physical hosts. Section 4 discusses related work that already addresses some of these issues. Finally, the findings are summarised and discussed in Section 5.

2 Characterising host virtualization

To frame the area of host virtualization from an IT management perspective, we examine new features, the use-cases ensuing from these features, as well as the differences between virtualization alternatives and the shortcomings of management systems.

2.1 Features of host virtualization

While virtual hosts retain a multitude of properties of physical hosts, it is the characteristic features of virtual hosts that allow new scenarios, leading to management challenges.

1. *Multiple VMs on single physical host* The primary feature of host virtualization is the ability to execute multiple VMs co-located on the same physical host.
2. *Isolation of applications below/at guest OS level* The virtualization facility isolates VMs from the hardware, as well as from each other. At best, the guest OS executes under the assumption that it runs exclusively on a hardware it controls. However, the degree of isolation (both vertically and horizontally) is determined by the virtualization strategy employed.
3. *Different guest OS brands/variants* Virtualization facilities may support VM instances running different guest OSs, e.g. Linux- or Windows-based OS variants.
4. *Different host/guest architecture* Virtualization facilities may support guest OS instances (and applications) built for a different machine architecture.
5. *Archiving and replication of VM* VMs are commonly prepared as images in a specific *image format* that contains the guest OS, drivers and application software. An image may be started and stopped, as one would a traditional OS on a physical machine, or they can be copied/replicated onto other physical hosts. Thus, an OS installation can be subject to the same operations as normal data files.
6. *Migration of VM* VMs can be transmitted over the network either in a frozen/stopped state, or while running. Thus, the physical location of a VM as well as its topological location within the network become volatile.
7. *Additional layer of software* Whichever the type of virtualization facility used, it introduces a distinct new layer of system software (II, in Figure 2).
8. *Nesting of virtualization facilities* A virtualization facility may be run within a VM, thus allowing recursive application of the abstraction provided by virtualization facilities.

2.2 Management scenarios

The delivery of host virtualization is driven by new scenarios that exploit the features identified as characteristic for virtual hosts. We exemplify typical scenarios in short to highlight their management requirements.

Consolidation The separation of applications onto different platforms (that is often dictated by the continued operation of "legacy" software) does benefit from the ability to operate different guest environments side by side, on the same physical hardware. Traditional "stove-pipe" architectures of OSS are being replaced by an increasingly horizontal distribution, i.e. looser coupling, and a larger numers of smaller packages.

The flexibility gained with respect to the selection of a deployment scheme does, however, quickly become a liability due to the lack of concepts of tools for determining and expressing business as well as operational constraints. The prerequisite, a reliable means of determining the performance of a specific deployment, is lacking.

Power management The quest to lower the consumption of electric power needed to operate computing equipment and the corresponding cooling facilities is driven by the rising energy prices and by environmental considerations. Host virtualization allows the consolidation of services running on underutilised physical machines by running them on virtual hosts instead. Consequently, a smaller number of machines are needed to provide the same set of services.

However, services are not accessed uniformly over time, and a machine providing service components may be overloaded by a burst of access to one of the services. Overprovisioning to absorb such bursts lowers or obviates the benefit of consolidating in the first place. Hence, mechanisms for dynamic tradeoff between energy consumption and processing power are required that support adherence to service level agreements (SLA), as well as concerns for availability in the event of a hardware failure that, in a consolidated environment, will impact a larger number of services.

High availability and load balancing In large-scale IT environments, cloud computing settings or grids, the load balancing potential of host virtualization together with the options to realise high-availability and manage maintenance windows constitute attractive features. Service availability can be assured in the event of maintenance windows, hardware failures or catastrophic incidents by means of VM replication and migration. Using the same mechanisms, load balancing within and between data centres can be achieved, moving heavily utilised VMs onto less loaded hardware.

We still lack reliable means to quantify and predict the load caused by single VMs and the cost for their migration. Additionally, the impact on availability, performance and security caused by re-placing VMs should be quantifiable at runtime, causing the need for adequate metrics.

2.3 Management systems

As suggested in Section 1, different techniques are employed to realise the basic virtualization function. Each exhibits different properties regarding performance, isolation of VMs against each other, abstraction from concrete machine architecture, abstraction from brand

and type of guest OS. Different data formats are used for the representation of VMs, and different resource management capabilities are implemented at dissimilar management interfaces.

Technical prerequisites for management have been developed by virtualization product vendors, ranging from the resource management functions for replication, archival, migration to higher-level schemes for HA applications or load distribution. They are offered as part of proprietary resource management tools. In order to leverage these functions within an integrated management, they must be included seamlessly in the management of networked systems, in management processes and, finally, become part of integrated management systems.

In general, every host virtualization product is supported by "its own", specialized resource management system, e.g. Citrix's XenCenter, VMWare's VirtualCenter, SWSoft's Parallels Management Console (for Virtuozzo) and so on. Support for multiple products in one management system is rare, for obvious commercial reasons. Producers of integrated management systems provide a loose coupling to (typically one of) the aforementioned resource management tools (in preference to integrating the management of virtual machine into their systems) although some management APIs have been published (e.g. [vmw07]) and standards for certain aspects of host virtualization (e.g. for VM image formats, [dsp09]) have been made available.

3 Classification of management issues

The management issues ensuing from virtualization technology can be organized according to existing conceptual frameworks such as the information, communication, functional, and organisational submodels (see e.g. [HAN99]) specified for management architectures, the reference processes being adopted for the process-oriented management of services (e.g. from the IT Infrastructure Library, ITIL or from the enhanced Telecom Operations Map, eTOM), as well as the consideration of the life-cycles of resources and services.

In this paper, we will constrain analysis to functional aspects at a systems management and IT management process level. We employ the classic functional areas (FCAPS) to provide an overall structure and include related aspects of process-oriented management.

3.1 Faults, incidents and problems

The handling of faults and errors constitutes an elementary management function addressed by the OSI functional area of the same name (technically) and by incident and problem management from a process-oriented point of view.
F-1 Symptom detection and incident handling. The symptoms experienced by users depend on the combination of possibly different guest OS instances running on a certain virtualization facility type on a given machine architecture.
F-2 Fault detection and localisation. VMs are no-longer bound to physical location.

Thus, after detecting a fault (and localising it in a logical topology) it must be localised in a (no-longer implicit) physical topology. Automated mechanisms are required to perform the mapping in real-time, over different virtualization facility types.
F-3 Root-cause analysis. An additional software layer (II) must be taken into account. Correctly determining the root cause in virtualized environments require aggregation, normalisation and correlation of monitoring data in layers I–III of the architecture sketched in Figure 2.
F-4 Fault recovery. Some recovery procedures common with physical machines (e.g. simply rebooting a recalcitrant server) may cause collateral damage in a virtualized environment. In addition, different strategies will be needed for the variable types of guest OS instances being run on a physical machine. Solutions based on automated planning of less invasive recovery steps may provide a safer means to restore services. *Reactive problem management* in the course of an escalation procedure (e.g. as suggested by the reference process specified by the ITIL) must deal with a higher complexity, both regarding the relevant software layers (II-IV in Figure 2), the actual VM instances being run and their interdependencies.
F-5 Risk analysis. The failure of a VM is implied by the failure of the physical machine it runs on. Where services are not protected by high-availability solutions, the distribution of VMs onto the base of physical hosts should ensure that the failure of any one host will impact the least number of services.
F-6 Proactive problem management. New VMs are often created by copying a VMs image; problem solutions for one VM should therefore be transferrable to those derived (copied) from it, preferably before incidents occur.
F-7 Knowledge database. ITIL's knowledge database for resolution processes must take into account the more flexible binding of services to (virtual and physical) hardware. Hence, existing schemas for *known problems* and *work-arounds* must be augmented accordingly.

3.2 Configuration and life-cycles

Configuration management is a function found in technical management of resources and services, frameworks for process-oriented management and, in an implicit manner, in the life-cycles of IT entities. The following challenges pertain to one or more of these categories.
C-1 Release management. The technique used to achieve the basic virtualization function poses requirements on the permitted combinations of software packages, including in particular guest OS brand and version. Hence, as known from e.g. critical system libraries, changes to one architectural layer (see Figure 2) may imply change to another. Releases must take into account the tuple (Virtualization facility, Operating System, Application) in dependence of the virtualization type. Versioning schemes (i.e. naming of versions) may have to be adapted to reflect this requirement. The software providing the virtualization function should be included in the Definitive Software Library proposed in the ITIL.
C-2 Configuration sets. Configuration parameters set in the virtualization facility may

constrain the configuration of guest OS instances. Consider e.g. the need for a properly configured physical network adapter under the virtualization facility's control, when executing VMs that rely on network connectivity. Automated distribution of (e.g. updated) configuration sets should take into account the same issues. Recursive application of the virtualization function perpetuates this issue to (upper) guest OS instances.

C-3 Normalisation. Guest OSs of different types will expose different configuration interfaces. To integrate configuration management across different OS and virtualization facility vendors and products in an automated manner, a common format and interface for executing configuration is necessary.

C-4 Mobile configurations. Migrating VMs may have to be re-configured, e.g when traveling to a different subnet. Therefore, configuration needs to become aware of a VM's *context*, as known from nomadic systems.

C-5 Audits. Configuration management audits are to ensure that only approved hardware/software is operated throughout an IT organisation's infrastructure. New issues introduced by host virtualization include if only permitted VM instances are being run and if such VMs are at a permitted location. Audits should therefore support a *zoning* concept with respect to management domains, security constraints etc.

C-6 Configuration management database (CMDB). To provide support for the process extensions discussed in this section, new CIs for VMs and virtualization facilities must be included in the CMDB. Hence, identity of the physical machine, the virtualization facility instance, as well as for guest OS instances must be recorded. As the VMs' location attribute's value is volatile, a means should be found to keep it current, in particular in environments that make use of the VM migration function.

C-7 Different life-cycle phases. Physical resources' life-cycles possess phases that are easily distinguishable. The same phases become blurred when dealing with virtual resources. For example, the provisioning phase may include the creation of an image, after which resources can be multiplied simply by instantiating a new component and applying a (new, adapted) configuration.

3.3 Accounting

Accounting of service usage (including the usage of VMs themselves, as a service) faces a number of challenges originating from the level of indirection introduced by virtualization technology.

A-1 Usage metrics. As the performance of a virtual system is no-longer strongly correlated with that of a physical machine, metrics based on load or volume must be revised when applying them to customers' usage.

A-2 Tracking. Automated changes (e.g. migration of VMs) performed according to higher-level objectives (e.g. availability or perfomance goals) should appear transparent to the management system.

A-3 Charging. Usage accounting and subsequent charging must be able to track the location of VMs pertaining to a given customer. Hence, either up-to-date location information is necessary, or the accounting/charging systems (e.g. usage metering agents, in

particular) must be resilient to such changes.
A-4 Heteroneous measurement points. Any usage accounting will be dependent on the properties (e.g. interfaces) of the physical machine as well as the OS available to a customer. To avoid constraints on VM placement, metering and accounting should yield consistent results across different architectures and OSs.

3.4 Performance, Availability and Capacity

The perceived performance of a (physical) machine executing multiple VMs is dependent on the number of VMs being executed, the load of these VMs and the overhead introduced by the virtualization facility.
P-1 Performance metrics and guarantees. Current virtual resource management systems provide means to limit and control VMs' use of single physical resources like CPU and memory. How to quantitatively gauge performance for single VMs in order to offer Quality of Service guarantees remains an open question. It is influenced by different architectures, OSs, VM numbers, nesting depth and application requirements.
P-2 Degree of efficiency. The degree of efficiency of virtualized environments would denote a metric for the amount processing power provided to VM users when compared to the "raw" amount provided by the physical infrastructure elements. Such a metric would allow for more consistent planning of capacity in the long term, as well as means for optimal loading of physical hosts.
P-3 Availability management. Availability management is a service management process and as such, it employs a *service* view on availability. Given that a service may be provided by several VMs whose location and co-location environment may vary, the level of its availability will have to be assessed in dependence of the VM deployment scheme.
P-4 Migration control. Migration of VM instances creates high network load, which may obviate the benefit from a migration motivated by performance. Concepts and algorithms are necessary that take into account the cost and benefit of migration, and that prevent "migration thrashing".
P-5 Mobile resource allocation. System resources (CPU, memory, bus) are allocated to VMs by current resource management systems, with respect to management policy derived from e.g. SLAs. Such allocations should be portable across physical hosts; relative values (e.g. 20% CPU time) should be adjustable according to the situation on a destination machine after migration.
P-6 I/O Management. Taking into account application I/O, as well as the I/O load incurred by loading VMs and by migration, a holistic method to control and allocate bus and network resources is necessary.

3.5 Security

Virtualization technology exposes new vulnerabilities of the confidentiality and integrity of data processed within VMs. Attacks on VMs (layer III) can originate from within the same layer, from an attacker with (partial) control of the virtualization facility (layer II), as well as along the known avenues of attack. In addition, organisational issues must be taken into account, a number of which have been sketched in [DHLgF08].

S-1 Conflicts of interest. Placing multiple customers' VMs on same physical machine can create potential conflicts of interest if the customers are competitors in the same field. Hence, change management constraints must be conceived and applied to the placements of VMs.

S-2 Isolation control. Security risk must be assessed regarding the guest OS type and the type of the virtualization facility. A metric for the isolation level is necessary, as well as alternatives to ensure a certain level of isolation.

S-3 Role management and access control. Environments where VMs are placed in customers' care, several administrative levels are necessary: the role managing the hardware and the virtualization facility, the role managing a set of VM instances, as well as managers of specific applications (e.g. DBMS or web servers) running within the VMs. A fine-grained role concept needs to be developed in order to reflect these requirements.

S-4 Patch management. Given the attack patterns mentioned above, maintenance of both virtualization facility and guest OSs is paramount (i.e. application of security patches). Multiple machine architectures, guest OS types aggravate the situation. Archived (i.e. not used) VM images must be taken into account.

S-5 Security zones. While migration of VMs is a fancy and useful feature, the technology itself does not take into account the security requirements of specific VMs. As a consequence, VMs may migrate between subnets with different levels of threat. Hence, mechanisms are required that pose additional constraints on the migration of VMs.

S-6 Security of images. While inactive or during migration, VM images could be maliciously modified. Since images do change while running, static signatures (as used in software distribution) do not suffice. Instead, a dynamic method is necessary to detect illegal modifications.

4 Related work

Scientific work in the domain of host virtualization is focused on certain topics of interest. Host virtualization is primarily being viewed as a *means* for IT management [McL08]. Many publications focus on provisioning, scheduling or distribution of VMs.

Provisioning and runtime overhead for VMs are found to be significant in Grid environments [SKF06], enough to warrant separate scheduling of VM provisioning to Grid users as well as caching strategies for VM images. Efficient installation of virtual clusters in high-performance computing environments (including the application of user-specific configurations to the VMs) is explored in [NMM07], and an installation system prototype is presented. In contrast, Qian et al. propose the use of "ghost" VMs, that are active but do not provide service, in order to quickly react to changing levels of demand in utility computing environments [QMZ$^+$07]. Kangarlou et al. demonstrate a method to record the state of distributed systems composed of virtual hosts in order to support *fault-recovery* by replaying the virtual environment [KXRE07]. Strategies for VM placement and migration are examined in [GIYC06] to provide brokarage of resource “slivers” between host servers, while [RFN07] describes an algorithm for distributing CPU time to VMs according to specified service levels. The concept of virtual workspaces is introduced in [KFF+05] to allow user-oriented provisioning of VMs in Grids, while [Beg06] and [MC07] propose language-based tools developed to provision virtual network components.

Koh et al. demonstrate *performance* interference effects in applications running on VMs [KKB+07] in spite of the toted isolation of VMs against each other. Performance of VMs running on multicore clusters have been analysed using synthetic benchmarks [RKGS08], and recommend to place application components in different VMs in order to avoid interference effects between the schedulers on layers II and III.

A scheme for policy-based *power management* in large-scale environments is proposed in [NS07], taking into account the power management capabilities built into common hardware (i.e. CPU). An approach for managing the tradeoff between power saving and performance is proposed in [SWHK08] that, while not specifically aimed at virtual hosts, does take into account the service levels assured to customers.

Security management challenges introduced by virtualization are discussed in [DHLgF08], taking into account technical issues as well as process-related issues. An architecture for access control in virtualized environments is presented in [PSC+07], which proposes a layered approach to access control that differentiates between access to the layers II–IV shown in Figure 2.

As suggested by the challenges presented in Section 3, the formalisation of data regarding VMs is of high importance. Current versions of the Common Information Model (CIM) [dsp07b] family of standards include a profile for virtualization containing classes for the representation of virtual environments [dsp07c, dsp07a] . Based on that profile, [JCKL07] proposes an approach for policy-based management for virtual resources modeled with the CIM. For iSCSI, a protocol often used for providing remote storage facilities to VMs, a management information base (MIB) module has been standardised [BKMM06], though it does not take into account the users (i.e. VMs) of the storage devices.

5 Discussion and conclusions

Host virtualization is a disruptive technology with respect to IT management as it negates ingrained assumptions and poses novel challenges in resource and service management. In particular, as communications services are dependent on effective OSS/BSS deployment, they are affected by the introduction of virtualization, and dependent on its efficient and effective operation. We have classified management issues and current work according to virtualization features and functional areas in Sections 3 and 4 and obtained a map of virtualization management challenges and addressed topics.

Figure 4 summarises the management challenges formulated in this section in a matrix of the characteristic features identified in Section 2 (horizontally) and the functional areas (vertically). Every requirement is entered in each relevant cell. The count of management issues is noted below and at the right of the diagram. The frames around some cell groups indicate relevant literature that has addressed specific feature/area pairs is noted in the matrix and discussed, along with relevant standardisation documents, in Section 4.

It is notable that research efforts appear concentrated in some management domains, e.g. in performance management. In other areas, e.g. accounting, a careful search has failed to identify relevant conceptual work, though the topic has been addressed in technical manuals (e.g. [AAT04]) associated with virtualization products. Even though no mere selection of related work can be said to represent a research domain faithfully, it is apparent that the more thoroughly addressed topics are, in fact, those that have been identified as "virtualization issues" early on; those not having been addressed yet (again: in this selection) are those that today pose difficulties in large-scale deployment of host virtualization, e.g. in cloud computing settings.

	Multiple VMs on physical host	Isolation	Different machine arch.	Different guest OSs	Archiving and replication	Migration	Additional software layer	Nesting of virt. facilities	
Fault	F-2 F-5		F-1 F-4 F-7	F-1 F-4 F-7	F-6 [KXRE07]	F-2 F-6 F-7	F-4 F-3		14
Configuration	C-5 C-6	C-2	C-1 C-3 [KFF+05]	C-1 C-3	C-7 [NMM07] [Beg06] [MC07]	C-4 C-5 C-6	C-1	C-2	13
Accounting	A-1 A-3	A-1	A-4	A-4		A-3 A-2			7
Performance	P-1 P-2 P-3 [QMZ+07] [SWHK08] [BFN07]	P-1 P-2 [RKGS08] [KKB+07]	[NS07] P-1 P-2	P-1 P-2	P-3 P-6 [SKF06]	[GIYC06] P-4 P-3 P-5	P-6	P-1 P-2	17
Security	S-1 S-3	S-2 [DHLgF08]	S-2 S-4	S-2 S-4	S-4 S-6	S-1 S-5 S-6	S-3 [PSC+07]	S-3	14
	11	5	10	10	10	14	5	4	

Figure 4: A map of host virtualization management

Taking into account the trend towards virtual network components, it appears that presently we will be confronted with managing whole virtual infrastructures residing on single physical hosts (i.e. virtual network elements and VMs). The management issues discussed in this paper indicate that managing virtual hosts across functional areas requires additional and current management information. The definition of a management information base (MIB) for virtual hosts, as well as for virtualization facilities, i.e. a *Virtualization MIB* would constitute a step towards integration, and it can draw on existing modelling work [dsp07c, BKMM06].

5.1 Viable management paradigms

Several management paradigms seem related to the challenges described in this paper: in some cases, they offer avenues for resolving management issues.

Extensions to service management processes Though primarily aimed at a higher-level management of IT services, process-orientated management should be extended by adding VM-specific information items to process artefacts and augmenting activities accordingly. Several requirements and issues given in the paper demand such extensions, e.g. regarding fault localisation and resolution (issues *F-1*... *F-4*) relevant in incident and problem management with respect to services, performance guarantees (issue *P-1*) that must correspond to service level objectives, conflict of interest issues (*S-1*) that may have to be controlled to fulfil contractual obligations, just to name a few examples. The additional architectural layer introduced by host virtualization together with the flexibility with respect to VM placement and life-cycle management makes necessary the introduction of effective metrics in the domains of performance and security management.

Self-management Virtualization solutions already use self-management capabilities, e.g. for high-availability or resource allocation functions. As a means to eliminate management labour through automation, the introduction of additional self-management aspects offers an attractive path. In particular, taking into account, the instant potential horizontal scalability of virtual infrastructures, automation at the resource level seems a necessity, if effort in higher-order management disciplines (i.e. network and service management) is to be contained.

Policy-based management Many management requirements concern the specification of constraints and policies on different aspects of host virtualization (VM co-location, migration etc). Policy-based management (PbM) seems a sufficiently flexible means to cope with the additional complexity of virtualized environments by allowing a codification of operational requirements, e.g. when considering migration control (*P-4*) in the context of isolation requirements (*S-1, S-2*). Though the issue of policy conflicts has still not been solved conclusively, it is being introduced in practice in management systems for virtual resources. An extension to higher-order management, while taking into account the specifics of VMs, seems a natural step to be taken.

5.2 Outlook on future work

The classification of management issues and relevant work presented in this paper is being offered as a basis for discussion on one hand—as no such classification of an area under active research can be claimed to be exhaustive. On the other hand, it serves as the foundation for current and future research projects on the topic of management of virtualized infrastructures.

While this paper has addressed the domain of *host* virtualization in particular, virtualization in itself is found in a multitude of domains: virtualization of network segments (e.g. the provisioning of VLANs) and of the networks themselves (i.e. virtual private networks) has been used and studied for some time.

Host virtualization itself acts as a driver for the research and production of virtualization solutions for I/O devices, sometimes called *I/O virtualization*. The devices in question include host-bus adapters (HBAs) and network interfaces (NICs), developed in order to serve the higher I/O load placed on physical machines that support multiple VMs at the same time. Super-scalar hardware structures are being designed for this purpose, e.g. NICs and HBAs supporting multiple transmission channels that may be assigned to single VMs or VM groups. Obviously, in this case, components that were outside the scope of system management must be utilised in accordance to *external* management goals, e.g. a VM may be assigned a HBA channel of its own if service level agreements mandate it.

Thus, virtualization as a concept is being introduced in most aspects of computing and communications infrastructures—and even services are said to be "virtualized" in some cases. At the same time, existing management resources are not yet adapted to deal with the additional layer of indirection and abstraction that is being introduced.

A remedial approach that bypasses the problem of proprietary, heterogeneous management solutions is the creation of a framework for the management of virtualized infrastructures. Such a framework should include all virtualization domains (hosts, networks, I/O ...), encompass the most important technical and process-oriented management areas and disciplines and allow an integration of the management of virtualized infrastructures with existing management concepts and products. The work presented in this paper is intended as a step on the way to such a framework.

Acknowledgment

The author wishes to express his thanks to 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. D. Kranzlmüller and Prof. Dr. H.-G. Hegering is a group of researchers at Ludwig-Maximilians-Universität, Technische Universität München, the University of the Federal Armed Forces and the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences. http://www.mnm-team.org.

References

[AAT04] Erich Amrehn, Ronald Annuss, and Arwed Tschoeke. Accounting and monitoring for z/VM Linux guest machines. Red book, IBM, May 2004.

[Beg06] Kyrre Begnum. Managing large networks of virtual machines. In *LISA '06: Proceedings of the 20th conference on Large Installation System Administration Conference*, pages 16–16, Berkeley, CA, USA, December 2006. USENIX Association.

[BKMM06] Mark Bakke, Marjorie Krueger, Tom McSweeney, and James Muchow. Definitions of Managed Objects for Internet Small Computer System Interface (iSCSI). RFC 4544, IETF, May 2006.

[DHLgF08] V. Danciu, W. Hommel, T. Lindinger, and N. gentschen Felde. Adaptive defense measures against the security hazards induced by systems virtualisation. In *Proceedings of the 2008 Workshop of HP Software University Association (HP–SUA)*, Marrakech, Morocco, June 2008. Hewlett Packard, Infonomics Consulting.

[dsp07a] CIM System Virtualization Model White Paper. White Paper DSP2013, Distributed Management Task Force, November 2007.

[dsp07b] Common Information Model (CIM) Infrastructure. Specification DSP00004, Distributed Management Task Force, November 2007.

[dsp07c] System Virtualization Profile. Specification DSP1042, Distributed Management Task Force, August 2007.

[dsp09] Open Virtualization Format Specification. Specification DSP0234, Distributed Management Task Force, February 2009.

[GIYC06] Laura Grit, David Irwin, Aydan Yumerefendi, and Jeff Chase. Virtual Machine Hosting for Networked Clusters: Building the Foundations for "Autonomic" Orchestration. In *VTDC '06: Proceedings of the 2nd International Workshop on Virtualization Technology in Distributed Computing*, page 7, Washington, DC, USA, 2006. IEEE Computer Society.

[GKM+03] S. Graupner, R. König, V. Machiraju, J. Pruyne, A. Sahai, and A. van Moorsel. Impact of Virtualization on Management Systems. Geneva, Switzerland, July 2003. HP OpenView University Association, HPOVUA 2003.

[HAN99] H.-G. Hegering, S. Abeck, and B. Neumair. *Integrated Management of Networked Systems – Concepts, Architectures and their Operational Application*. Morgan Kaufmann Publishers, ISBN 1-55860-571-1, 1999.

[JCKL07] Neeraj Joshi, Seraphin Calo, David Kaminsky, and Jorge Lobo. CIM-SPL Policies and Virtualization. In *1st International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and New Technologies*, volume 2007, Toulouse, France, October 2007. DMTF, IRIT and Univerist Paul Sabatier.

[KFF+05] Katarzyna Keahey, Ian T. Foster, Timothy Freeman, Xuehai Zhang, and Daniel Galron. Virtual Workspaces in the Grid. In José C. Cunha and Pedro D. Medeiros, editors, *Euro-Par*, volume 3648 of *Lecture Notes in Computer Science*, pages 421–431. Springer, 2005.

[KKB+07] Younggyun Koh, R. Knauerhase, P. Brett, M. Bowman, Zhihua Wen, and C. Pu. An Analysis of Performance Interference Effects in Virtual Environments. *2007 IEEE International Symposium on Performance Analysis of Systems & Software*, 0:200–209, 2007.

[KXRE07] Adarlan Kangarlou, Dongyan Xu, Paul Ruth, and Patrick Eugster. Taking Snapshots of Virtual Networked Environments. In *Second International Workshop on Virtualization Technology in Distributed Computing (VTDC 2007)*, Reno, Nevada, USA, November 2007.

[MC07] Fermin Galán Márquez and David Fernández Cambronero. Distributed Virtualization Scenarios Using VNUML. In *1st International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and New Technologies*, volume 2007, Toulouse, France, October 2007. DMTF, IRIT and Universit Paul Sabatier.

[McL08] Laurianne McLaughlin. Virtualization in the Enterprise Survey: Your Virtualized State in 2008. CIO Research, CXO Media Inc., January 2008.

[NC05] S. Nanda and T. Chiueh. A survey on virtualization technologies. Technical Report TR-179, Department of Computer Science, State University of New York, February 2005.

[NMM07] Hideo Nishimura, Naoya Maruyama, and Satoshi Matsuoka. Virtual Clusters on the Fly - Fast, Scalable, and Flexible Installation. In *CCGRID '07: Proceedings of the Seventh IEEE International Symposium on Cluster Computing and the Grid*, pages 549–556, Washington, DC, USA, 2007. IEEE Computer Society.

[NS07] Ripal Nathuji and Karsten Schwan. VirtualPower: Coordinated Power Management in Virtualized Enterprise Systems. In *Proceedings of Symposium on Operating Systems Principles (SOSP'07)*, Stevenson, Washington, USA, October 2007.

[PSC+07] Bryan D. Payne, Reiner Sailer, Ramón Cáceres, Ron Perez, and Wenke Lee. A layered approach to simplified access control in virtualized systems. *SIGOPS Oper. Syst. Rev.*, 41(4):12–19, 2007.

[QMZ+07] Hangwei Qian, Elliot Miller, Wei Zhang, Michael Rabinovich, and Craig E. Wills. Agility in Virtualized Utility Computing. In *Second International Workshop on Virtualization Technology in Distributed Computing (VTDC 2007)*, Reno, Nevada, USA, November 2007.

[RFN07] Fernando Rodríguez, Felix Freitag, and Leandro Navaro. Towards intelligent management in VM-based resource providers. In *1st International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and New Technologies*, volume 2007, Toulouse, France, October 2007. DMTF, IRIT and Universit Paul Sabatier.

[RKGS08] Adit Ranadive, Mukil Kesavan, Ada Gavrilovka, and Karsten Schwan. Performance Implications of Virtualizing Multicore Cluster Machines. In *2nd Workshop on System-level Virtualization for High Performance Computing (HPCVirt 2008)*, Glasgow, Scotland, March 2008.

[SKF06] Borja Sotomayor, Kate Keahey, and Ian Foster. Overhead Matters: A Model for Virtual Resource Management. In *VTDC '06: Proceedings of the 2nd International Workshop on Virtualization Technology in Distributed Computing*, Washington, DC, USA, 2006. IEEE Computer Society.

[SN05] James E. Smith and Ravi Nair. The Architecture of Virtual Machines. *Computer*, 38(5):32–38, 2005.

[SWHK08] Malgorzata Steinder, Ian Whalley, James E. Hanson, and Jeffrey O. Kephart. Coordinated Management of Power Usage and Runtime Performance. In *Pervasive Management for Ubiquitous Network and Services. Proceedings of the IEEE/IFIP Network Operations & Management Symposium*, Salvador, Bahia, Brasil, April 2008.

[vmw07] *Programming Guide. VMware Infrastructure SDK 2.5.* VMware, Inc., Palo Alto, CA, USA., November 2007.

[Wol07] Chris Wolf. Let's Get Virtual: A Look at Today's Server Virtualization Architectures. In-depth research overview, Burton Group, May 2007.