Evaluation of an Accounting Model for Dynamic Virtual Organizations

Martin Waldburger · **Matthias Göhner** · **Helmut Reiser** · **Gabi Dreo Rodosek** · **Burkhard Stiller**

Abstract Accounting of Grid resource and service usage determines the central support activity for Grid systems to be adopted as a means for service-oriented computing in Dynamic Virtual Organizations (DVO). An all-embracing study of existing Grid accounting systems has revealed that these approaches focus primarily on technical precision, while they lack a foundation of appropriate economic accounting principles and the support for multi-provider scenarios or virtualization concepts. Consequently, a new, flexible, resource-based accounting model for DVOs was developed, combining technical and economic accounting by means of Activity-based Costing (ABC).

Driven by a functional evaluation, this paper pursues a full-fledged evaluation of the new, generically applicable Grid accounting model. This is done for the specific environment of the Leibniz Supercomputing Centre (LRZ) in Garching, Germany. Thus, a detailed evaluation methodology and evaluation environment is outlined, leading to actual model-based cost calculations for a defined set of considered Grid services. The results gained are analyzed and respective conclusions on model applicability, optimizations, and further extensions are drawn.

M. Waldburger, B. Stiller
University of Zurich, Department of Informatics (IFI), CH–8050 Zürich, Switzerland
E-mail: waldburger@ifi.uzh.ch, stiller@ifi.uzh.ch

M. Göhner, G. Dreo Rodosek
Universität der Bundeswehr München (UniBwM), D–85577 Neubiberg, Germany
E-mail: matthias.goehner@unibw.de, gabi.dreo@unibw.de

H. Reiser
Leibniz Supercomputing Centre, D–85748 Garching near Munich, Germany
E-mail: reiser@lrz.de

B. Stiller
ETH Zürich, Computer Engineering and Networks Lab (TIK), CH–8092 Zürich, Switzerland
E-mail: stiller@tik.ee.ethz.ch

Keywords Grid Computing · Accounting · Dynamic Virtual Organization · Activity-based costing

1 Introduction

Grid service accounting constitutes a central functional support activity in both, research-oriented and business Grid systems, as it facilitates the creation of service and resource usage records. Accounting relies on successful user authentication and authorization. Once access to a resource or respectively a service is granted, resource usage has to be accounted reliably. This results from the fact that accounting data becomes retrievable for auditing purposes or—in a fully competitive environment—it is finally transferred into charging records which in turn will be equipped by monetary values so that a bill to the service consumer can be issued. These steps are reflected by AAA (Authentication, Authorization, Accounting) [19][27] and its extended view, A4C (AAA plus Auditing and Charging) [10][20].

Accounting for Grid systems represents an important research focus, since it constitutes on the one hand the key mechanism for commercial electronic services to be offered and charged to customers and on the other hand, accounting data potentially contain valuable information for a Grid service provider regarding current and past service usage as well as resource consumption. Such information can be used for charging purposes as well as for internal optimization processes or service portfolio optimization. Both require accountable units that equip a service provider with significant information that correlates closely with chosen optimization criteria. For instance, a service provider may want to optimize its cost-benefit ratio. For that purpose Grid service accounting is required to produce records that allow this service provider to identify and classify the relevant set of cost drivers.

In the same way as Grid service accounting is of key importance to outlined reasons like successful commercialization and cost management, the respective steps of Grid service accounting have to build on a solid theoretical basis being represented by the appropriate underlying Grid accounting model. This Grid accounting model is required to satisfy multiple demands. These comprise technical requirements such as precision and scalability in obtaining accounting records, and, equally important, economic requirements such as a sound support of established cost accounting methods from the accounting across organizational boundaries in DVOs.

There are many accounting approaches for Grid systems available, which lack a sound economic accounting basis as they are highly specific to the considered application case so that they are not generically applicable [15]. To overcome these shortcomings, a resource-driven and activity-based accounting model for DVOs—as implemented by Grid systems—was developed [15][17]. The generic model which is described in greater detail in Section 2.3 is used to calculate costs incurred for a given Grid service in the context of a DVO. The developed model has proven to be a highly promising approach from a functional point of view [15].

Based on the existing conceptual evaluation of our presented approach in [15], a full-fledged assessment of this model in existing Grid environments needs to be undertaken. This evaluation constitutes the main focus of this work. It is done by applying the generic model to the Grid infrastructure operated by the LRZ, the Leibniz Supercomputing Centre in Garching near Munich, Germany [25]. The evaluation's main goal consists in applying the conceptually evaluated Grid accounting model to an existing operational Grid infrastructure in order to reveal the key set of practical aspects relevant for model application and to determine model improvements and extensions. In particular, the model is assessed by means of three dimensions. In consideration of the model's overall aim to calculate costs of a Grid service, the evaluation addresses achieved model functionality, available and used means of model parametrization, and serviceability regarding the respective LRZ application context.

Accordingly, the remainder of this paper is structured as outlined in Figure 1. Section 2 provides an overview of related work for accounting in DVOs. Driven by the analysis of existing Grid accounting approaches (Section 2.1) and the derived requirements on Grid accounting (mentioned explicitly in Section 2.3), this includes in particular a presentation of the respective key characteristics of previous achievements, namely the developed DVO service model (Section 2.2), a comprehensive Grid resource classification (Section 2.4), and the developed Grid accounting model for DVOs (Section 2.3).

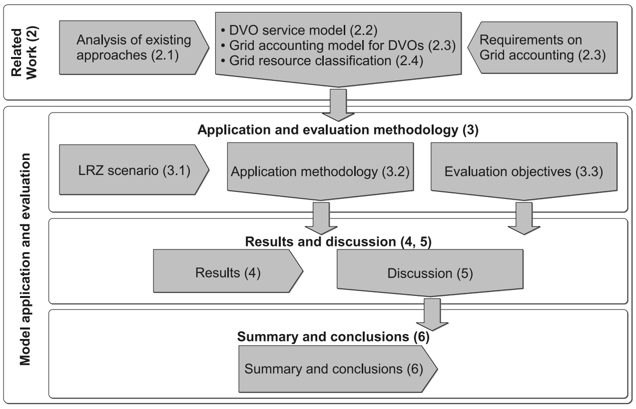

Fig. 1 Paper Structure (Sections in Brackets)

Later sections address this work's core focus determined as the application and evaluation of the generic Grid accounting model to the LRZ environment. This builds on a detailed description of the used application and evaluation methodology in Section 3, covering an in-depth investigation of the considered LRZ Grid infrastructure and the respective multi-domain Grid accounting scenario (Section 3.1), an all-embracing description on necessary steps to apply the Grid accounting model to the determined scenario and LRZ infrastructure (Section 3.2), and a definition of objectives and requirements for model application assessment. According to those outlined application and evaluation methods, model calculations and the according results are presented in Section 4 and discussed in Section 5. Driven by the gained insights, the work is summarized and the respective conclusions are drawn in Section 6, including proposed adaptations of the Grid accounting model.

2 Related Work

In this section, related work addressing the research domain of Grid accounting is presented and relevant concepts are discussed. Herefore, Section 2.1 contains an overview of existing Grid accounting approaches which are evaluated against a list of 23 identified criteria, which have been derived on comprehensive requirements analysis as well as various accounting-specific use cases. Moreover, as a sound theoretical basis for successful model application and evaluation, terminology in use and those key mechanisms for Grid service accounting in DVOs need to be outlined. This covers in particular the inspection of core achievements from previous work, namely those developed core models—DVO service (Section 2.2) and Grid accounting model (Section 2.3)—as well as an all-embracing classification of Grid resources and possible accountable units (Section 2.4).

2.1 Overview and Evaluation of Existing Accounting Systems

Based on a comprehensive survey on Grid accounting approaches in [11] and [15], the following provides an overview of existing accounting systems and tools from European as well as international Grid projects and finally presents an evaluation of fundamental characteristics as shown in Table 1. In the survey, the following accounting systems were analyzed:

- Accounting processor for Event Logs (APEL) [7]
- Distributed Grid Accounting System (DGAS) [2]
- Grid Accounting Services Architecture (GASA)/GridBank [4]
- Grid Based Application Service Provision (GRASP) [16]
- Grid Service Accounting Extensions (GSAX) [5]
- Multi-organisation Grid Accounting System (MOGAS) [26]
- Nimrod/G [6][3]
- SweGrid Accounting System (SGAS) [29]

In consideration of technical aspects, Table 1 depicts that, by focusing only on the accounting of physically existing Grid resources, none of the examined approaches addresses a concept for service and resource virtualization. Additionally, existing systems do not provide mechanisms for the accounting of composed virtual services and virtual resources as they are usually offered within multi-provider Grid environments. These are both key requirements for service provisioning and the according accounting in DVOs. Additionally, to some extent, only static environments with Grid resources of homogeneous nature and few accounting units are supported. Dynamic Grid environments with a high level of heterogeneity regarding services and resources, operating systems, and Grid middleware solutions are in most cases not taken into consideration.

Table 1 Evaluation of Existing Systems (+ "Yes", (+) "In parts", – "No", n.s "Not Specified") [15]

Criteria	Accounting System							
	APEL	DGAS	GASA	GRASP	GSAX	MOGAS	Nimrod/G	SGAS
Interoperability and portability	(+)	(+)	(+)	n.s.	(+)	(+)	+	+
Scalability	+	(+)	–	n.s.	+	(+)	+	+
Integration	(+)	(+)	(+)	n.s.	(+)	+	+	+
Inter-organizational accounting	+	+	+	n.s.	+	n.s.	n.s.	+
Flexibility and extensibility	+	n.s.	+	n.s.	+	(+)	(+)	+
Support of existing standards	–	–	(+)	(+)	(+)	n.s.	n.s.	+
Support of multi-provider scenarios	–	–	–	–	–	–	–	–
Visualization of accounting data	+	–	–	n.s.	n.s.	+	n.s.	–
User transparency	n.s.	n.s.	n.s.	n.s.	n.s.	(+)	n.s.	(+)
Accounting of heterogeneous resources	(+)	+	+	(+)	n.s.	(+)	n.s.	–
Accounting of virtual resources	–	–	–	–	–	–	–	–
Accounting of virtual services	–	–	–	–	–	–	–	–
Virtualization concept	–	–	–	–	–	–	–	–
Support of high dynamics	+	(+)	(+)	n.s.	n.s.	(+)	+	+
Security	n.s.	+	+	n.s.	+	+	n.s.	+
Standardized, generic interfaces	–	–	–	n.s.	(+)	n.s.	+	(+)
Support of various accountable units/metrics	+	+	+	n.s.	+	n.s.	n.s.	–
Precision and abundance	+	+	+	+	+	+	n.s.	+
Support of different accounting policies	+	+	n.s.	n.s.	+	-	n.s.	(+)
Reliability and fault tolerance	n.s.	n.s.	(+)	n.s.	n.s.	n.s.	n.s.	+
Administration and management	n.s.	(+)	n.s.	n.s.	n.s.	n.s.	n.s.	+
Verification	n.s.	+	+	n.s.	n.s.	n.s.	+	+
Open source	+	+	+	–	–	n.s.	+	+

Beside the examined Grid accounting systems and tools, [18] presents a high-level description of an infrastructure comprising accounting, banking as well as electronic payment services that are used for service-oriented Grid computing systems. This mainly theoretical approach only incorporates an accounting of elementary Grid services and physically existing Grid resources. Compound virtual Grid services and resources in multi-provider domains of dynamic Virtual Organizations are not taken into consideration. Additionally, the proposed architecture mainly focuses on payment issues and does not consider any aspects addressing the determination of costs incurred for a provided Grid service by combining technical and economic accounting, thus lacking an adequate economic basis.

In general, the study of existing approaches revealed that currently deployed Grid accounting systems mainly focus on technical precision and project-specific issues while they are not based on adequate economic cost accounting principles suitable for the accounting across organizational boundaries and DVOs. In addition, present accounting systems and tools usually have been developed for specific application areas comprising homogeneous hardware platforms and uniform technical infrastructures thus being not generically applicable on highly dynamic Grid environments [8][11]. Moreover, in many cases, the focus of existing accounting approaches is mainly on technical optimization criteria like measurement procedures and metering points with regard to the acquisition of accounting relevant data. Despite the fact that existing systems as for example SGAS, DGAS and GASA consider economic aspects, *e.g.*, payment schemes and bank services, business aspects of accounting regarding methods of cost calculation and cost accounting are not taken into account by any approach.

Since the above identified missing characteristics of existing Grid accounting approaches are of key relevance to a technically and economically sound multi-domain Grid accounting, the need to develop an appropriate Grid accounting model for DVOs became apparent. This led to major achievements in the suitable DVO service (cf. Section 2.2) and Grid accounting models (cf. Section 2.3) on one hand and in a classification of different Grid resource types on the other hand (cf. Section 2.4). These results of previous work constitute a solid theoretical basis for the Grid accounting model's application and evaluation.

2.2 DVO Service Model

In previous work [15], a comprehensive service model for DVOs was developed taking into account the concept of resource and service virtualization within multi-provider Grid environments. This service model which reflects the provider's perspective is structured into two separate layers, *i.e.*, a Virtual Organization (VO) layer and a layer of underlying real organizations (RO) providing an adequate basis with respect to appropriate structure descriptions and possible compositions of virtual services and virtual resources provisioned within the context of DVOs.

Figure 2 illustrates a formal representation of this service model comprising all relevant entities as for instance VOs and ROs along with their elements, *i.e.*, real services (S) and real resources (R) as well as virtual services (VS) and virtual resources (VR). Moreover, the UML notation of the service model reflects possible types of interactions between involved elements as for example utilization, composition as well as a mapping between VO and RO layers. A detailed overview of the service model along with a description of its elements and fundamental characteristics, as well as a presentation of concrete examples regarding resource and service provisioning within DVOs can be found in [11][15].

2.3 Grid Accounting Model for DVOs

Based on the service model for DVOs introduced in Section 2.2 and driven by the analysis of existing Grid accounting approaches (*cf*. Section 2.1), a generic accounting model was proposed [15][17] that allows for the accounting of complex, composed virtual services and virtual resources in

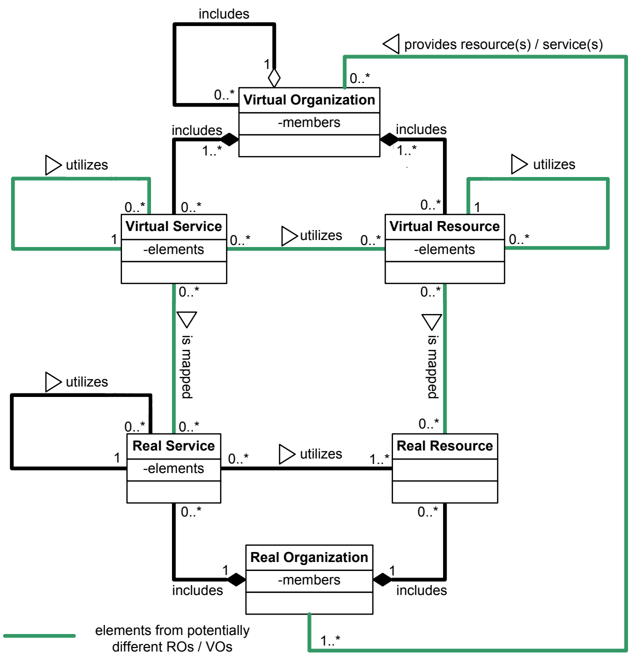

Fig. 2 Formal Representation of the Service Model [11]

multi-provider Grid environments, thus, going a step further than existing approaches.

The presented accounting model which focuses on economic and technical aspects was derived in accordance with a set of determined generic, DVO-specific requirements. Concrete examples are (i) compliance with the service model for DVOs, (ii) providing capabilities for bridging the concepts of cost accounting and technical accounting, (iii) support of various accountable units adequately reflecting resource consumption and service usage, as well as (iv) a high degree of flexibility, applicability, and extensibility for the use within highly dynamic Grid environments.

The proposed accounting model relies on two accounting concepts that are well-known in the domain of (economic) cost accounting: These are the Traditional Cost Accounting System (TCAS) and ABC [21][22]. TCAS relates to established, standard methods in economic cost accounting—also referred to as managerial or internal accounting. Hence, details on principles of TCAS can be found in text books on cost accounting, such as [23]. ABC is a widely accepted costing system that is particularly well suited for the accounting of electronic services [13]. In our Grid accounting model, TCAS and ABC are interconnected by means of so called service constituent parts, namely *Processing*, *Storage*, *Transferring*, and *Output*, representing a consistent set of building blocks every provisioned Grid service can be composed of. Figure 3 illustrates the fundamental idea of bridging the gap between TCAS and ABC by means of the identified service constituent parts along with their central role in the accounting process.

In addition, these four service constituent parts represent the basic hardware functionality within the context of Grid Computing, out of which any electronic service is assembled by some specific amount. The service constituent parts themselves are adapted to the specific resource they reflect. This is required, since typically different costs incur, when a job is run on different hardware or with specified service guarantees. Thus, in addition to interconnecting TCAS and ABC, these service constituent parts also interconnect economic and technical accounting. Technical accounting is defined as the "collection of resource consumption data for the purposes of capacity and trend analysis, cost allocation, auditing, and billing. Accounting management requires that resource consumption be measured, rated, assigned, and communicated between appropriate parties" [1]. Accordingly, the use of service constituent parts as a concept in order to configure activities for ABC links to the respective set of accountable units as needed for metering and accounting record preparation.

- **Processing** calculates costs for computation and data processing by using computational resources.
- **Storage** considers incurred costs for data storage and archiving by means of storage resources.

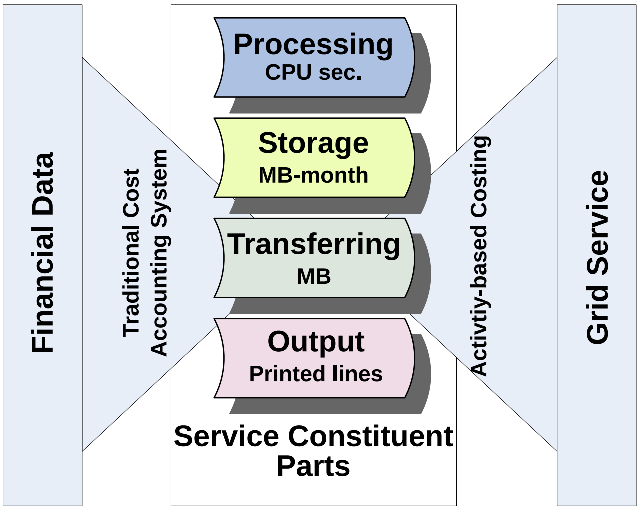

Fig. 3 Accountable Units Overview [15]

– **Transferring** reflects costs for transferring data within or between ROs or VOs respectively by use of network components.
– **Output** calculates costs for generated output, *e.g.*, printed documents, graphical representation of simulation results etc.

Moreover, in order to be able to allocate also other costs for service provisioning which are not chargeable to any of the above mentioned service constituent parts, a further generic service constituent part *Other* has been specified. Concrete examples for this service constituent part are organization-specific cost elements such as, *e.g.*, administrative cost that accrue due to service provisioning, but which cannot be mapped to a particular resource. Finally, the constituent part *External* is used to take costs into consideration that are associated with the usage of a service or a resource provisioned by an external provider as for example another VO. A detailed description of the identified service constituent parts along with concrete examples with respect to applicable metrics, relevant cost drivers, and associated costs can be found in [15].

These identified service constituent parts are resource-specific and mapped to activities. This means that the final IT product, *e.g.*, in form of a composed virtual service consists of a number of sub processes whereas sub processes are composed by activities, and activities are finally composed by service constituent parts serving as building blocks in the cost analysis process. In the example given in Figure 4, VO_1 offers a virtual service that is composed of two external services provided by RO_1 and RO_2. In addition to the costs incurred by sourcing those external services, additional costs as for instance for administrative activities are included on the VO level. Focusing on the first external service provided by RO_1, the example reflects the cost-relevant activities which are needed in order to provide this service to VO_1. Similarly, on level of RO_1, an external service is sourced from a third party, followed by RO_1's main process along with other cost elements that are not specified in greater detail at this stage. Within the administrative domain of RO_1, several steps that aggregate information are taken, leading in a top-down approach to a fine-granular process cost analysis, until, on the lowest level, the respective service constituent part assignment per real IT resource is conducted.

2.4 Grid Resource Classification

By means of those presented generic and extendable service constituent parts, our Grid accounting model provides the basis for a highly flexible, resource-based accounting in DVOs. In order to apply the model to a complex and heterogeneous environment such as the LRZ, however, an in-depth understanding of those resources of use in Grid systems is needed. Within the context of commercial and research-oriented Grid environments, *e.g.*, the D-Grid, a German-wide Grid infrastructure for establishing methods of e-Science in the German scientific community [9], a variety of different types of Grid resources having a high degree of heterogeneity can be identified. The basic requirement of the accounting system of supporting an accounting of various types of real as well as virtual Grid resources, which determine the basis for electronic service provisioning, implies the development of a taxonomy of Grid resources and possible sub types of resources.

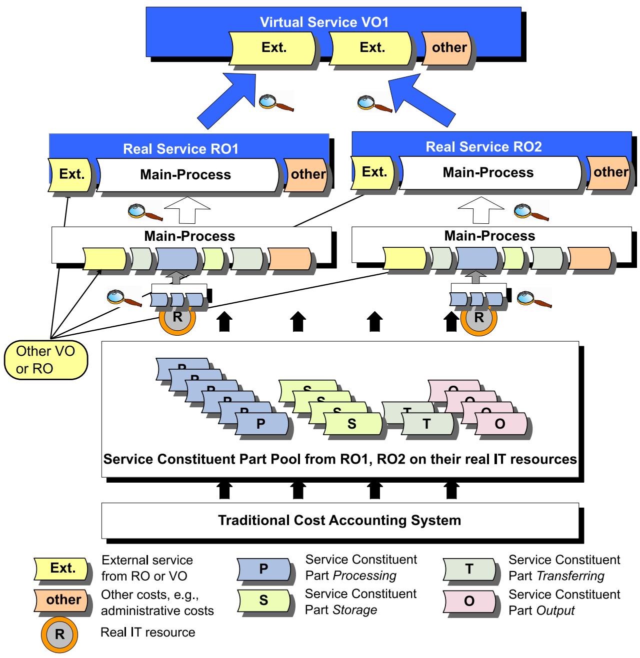

Fig. 4 ABC Accounting Model for DVOs

Group of resource	Possible sub groups	Examples	Possible accounting units
Computational elements	Multi processor systems	• Vector computer • Parallel computer • Cluster computer (e.g., IBM p690 Cluster etc.) • High-performance computer (e.g., SGI Altix 4700 etc.) • etc.	• CPU seconds • CPU hours • Wallclock time • Number of CPUs • Number of nodes • Number of computers/ computer systems • MIPS • etc.
	Single processor systems	• Desktop PCs (e.g., X86, x86_64, PowerPC etc.) • etc.	
	Hardware elements/ emulators	• FFT hardware (e.g., special hardware emulators etc.) • Co processor • etc.	
Storage elements	Primary storage elements	• Main memory • Caches (e.g., special harddisk caches etc.) • etc.	• Number of page accesses • Main memory (max.) • Main memory (avg.) • etc.
	Mass storage elements	• RAID systems • Tape systems • Archive systems • etc.	• Used storage (MB/GB/TB) • Used storage x time • etc.
Databases	Relational databases	• mySQL • Oracle • IBM DB2 • etc.	• Number of accesses • Utilization time • Value of extracted information • etc.
	XML databases	• eXist • Xindice • Tamino • etc.	
Network components	-	• Router • Switch • Gateway • Communication networks (e.g., LANs, WLANs, WANs) • etc.	• Bandwidth • Transferred data (MB/GB/TB) • etc.
Software components/ libraries	-	• Software licenses (e.g., medical software etc.) • Program libraries • Specialized software • etc.	Costs for: • Software licenses • Applications • Access to libraries • etc.
Resources for data aquisition	-	• Gauging station (e.g., virtual telescope, observatories) • Specialized hardware (e.g., electron microscopes, etc.) • Sensors • etc.	• Number of accesses • Utilization time • etc.
Further resources/ resource types	-	• Information systems • Visualization components • Administration/support • QoS parameters (e.g., priorities etc.) • etc.	Due to different characteristics of resources very application-specific accounting units

Fig. 5 Classification of Grid Resources and Possible Accounting Units

Therefore, a classification of different Grid resource types is presented. This classification provides an appropriate basis for the identification of accounting units and metrics adequately reflecting resource consumption and service usage. Basically, the following set of Grid resources can be identified:

– Computational elements
– Storage resources
– Network components
– Databases/information repositories
– Software components and licenses
– Specialized hardware and scientific devices

In Figure 5 a detailed classification of Grid resources and possible sub groups along with a list of appropriate accounting units per resource type is outlined, thus providing a useful basis for the specification of accounting units for the identified service constituent parts as described in Section 2.3.

3 Application and Evaluation Methodology

In accordance with service and accounting model characteristics, and in consideration of the described Grid resources, the used methodology for application and evaluation of the presented Grid accounting model needs to be outlined. Section 3.1 determines an LRZ-specific scenario for Grid accounting model application and evaluation. This involves detailed considerations of LRZ infrastructure and Grid services as well as an overview of financial, cost-related input data. While Section 3.2 outlines those functional steps required for Grid accounting model application, the set of relevant evaluation objectives and requirements is determined in Section 3.3.

3.1 LRZ Scenario Definition

The heterogeneous supercomputing infrastructure of the LRZ constitutes a complex application environment for the Grid accounting model at hand. Section 3.1.1 introduces the LRZ Grid infrastructure components. This is followed by presenting an elaborate accounting scenario in Section 3.1.2. The LRZ Grid infrastructure and the scenario provide the basic frame for subsequent model application—in particular with respect to cost calculations—and evaluation tasks.

3.1.1 LRZ Grid Infrastructure

As a service provider for scientific high performance computing, the LRZ operates computation systems for use by educational institutions in Munich, Bavaria as well as on a nationwide level. Beyond operation of system hardware, services offered at the LRZ also comprise backup/archive, Grid Computing as well as training courses on usage of HPC (High Performance Computing) systems, parallel programming and optimization [24].

The LRZ infrastructure encompasses several computing facilities. These consist, *e.g.*, of the new National Supercomputer “Höchstleistungsrechner in Bayern II” (HLRB II) based on SGI’s Altix 4700 platform which is optimized for high application performance and high memory bandwidth. Within the second phase of installation, the HLRB II has currently a total number of 9’728 CPU cores based on Intel Itanium2 Montecito Dual Core processors with an overall peak performance of 62.3 TFlop/s and 39 TByte of system memory as well as 600 TByte of direct attached disks. Current projects performed on the HLRB II reside in the domain of applied mathematics, astrophysics, biosciences, chemistry, and computational fluid dynamics etc. [24].

Moreover, the LRZ consists of several Linux-based cluster systems of varying size, performance, interconnect, and architecture (32 and 64 bit Intel processors) comprising close to 700 CPU cores in total. In 2008, the LRZ Linux clusters are extended to more than 3’500 CPU cores. The LRZ Linux clusters offer shared and distributed memory, varying available memory sizes, parallelization based on message passing (MPI), and shared memory parallelization.

The main focus of the Linux cluster systems is the development and testing of HPC applications as well as capacity computing.

The computing facilities offered at the LRZ—in particular the Linux clusters—are characterized by a high degree of heterogeneity with respect to underlying hardware platforms, numbers of processors, sizes of shared memory, and batch systems. In addition, three different kinds of Grid middleware solutions (Globus Toolkit [14], UNICORE [30] and gLite [12]) are currently in productive use resulting in a heterogeneous Grid infrastructure.

3.1.2 Multi-domain Grid Accounting Scenario

In the following, a fictitious scenario addressing the utilization and the accounting of a complex virtual service is presented in detail. This scenario can be seen as a concrete instantiation of the service model introduced in Section 2.2. It serves as a basis for the evaluation of the proposed accounting model. Moreover, the example scenario is enhanced with concrete values and parameter settings reflecting the usage of a compound virtual service consisting of several underlying services and resources which can be seen as building blocks the virtual service is composed of. Based on existing real-world accounting data reflecting service usage and resource consumption within the layer of the underlying real organizations, *i.e.*, the Grid infrastructure at the LRZ, an abstraction with regard to the virtual resources and virtual services provisioned within the layer of the Virtual Organizations is being performed.

This multi-domain scenario as depicted in Figure 6 comprises two VOs (VO_1 and VO_2) and two underlying ROs consisting of the LRZ which is part of VO_1 as well as a fictitious Grid service provider being part of VO_2 thus spanning multiple administrative domains. For reasons of simplification, the presented scenario only contains a 1:1 mapping between involved VOs and the underlying ROs, *i.e.*, one VO consists of exactly one RO. In real-world Grid environments, the normal case is that several ROs jointly participate in one or multiple VOs, respectively.

Within the considered example scenario, VO_1 offers a virtual simulation service (VS_1) performing large, three-dimensional simulations of turbulent flows and reactive flows in complex geometries. Accordingly, VS_1 comprises several data- and computation-intensive tasks. In the scenario, the simulation service VS_1 provisioned by VO_1 consists of several (sub) elements, *i.e.*, real as well as virtual services and resources which are offered by different organizations (VOs and ROs) jointly contributing the offered functionality of the virtual service VS_1.

The virtual simulation service VS_1 comprises a virtual computation service (VS_2) which is provided upon a compound virtual computation resource (VR_1) on which com-

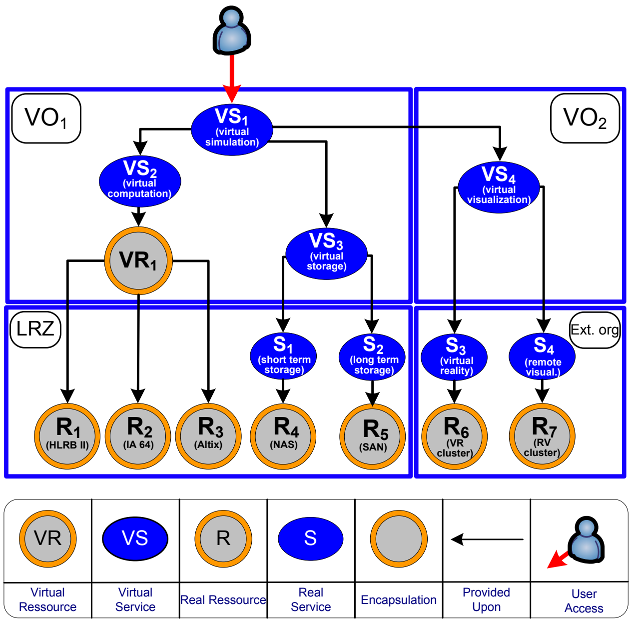

Fig. 6 Fictitious Accounting Scenario

plex calculations are performed. Moreover, VS_1 makes use of a virtual storage service (VS_3) being composed of two underlying real storage services (S_1 and S_2) offered within the LRZ. VS_3 is used for the archival storage of acquired simulation results. The real data services S_1 and S_2 which are responsible for the resource management coordination as well as the transparent storage of the data are provided upon physically existing storage resources R_4 and R_5. Finally, the virtual simulation service comprises a visualization service (VS_4) offered by an external provider (VO_2) in order to graphically illustrate the simulation results which are forwarded from the computation service VS_2.

Within the considered scenario, 19 percent (=512 processors) of the supercomputer HLRB II (R_1) are available for the execution of the user job. In addition, negotiated Quality-of-Service (QoS) parameters with respect to, *e.g.*, execution time of a user job have to be met. Therefore, besides the HLRB II also a part of the 64-Bit cluster IA 64 (R_2) of the LRZ infrastructure comprising a total of 220 processors as well as 25 percent (=32 processors) of the Linux Cluster based on the SGI Altix 3700 Bx2 (R_3) are used as part of the virtual computation resource VR_1. In order to perform the necessary calculations of the simulation service the 512 processors of the supercomputer HLRB II are used for 2.5 hours with a memory utilization of 2 GByte per processor whereas the physically existing resource R_2 is utilized for 4 hours along with a utilization of 1 GByte per processor of primary storage. Finally, 25 percent of the SGI Altix 3700 cluster is utilized for a time period of 6 hours together with a temporary consumption of main memory of 1.5 GByte per processor.

Simulation results with an overall size of 7 TByte are archived on storage resources at the LRZ by use of the two real data services S_1 and S_2. In this context, frequently used simulation results with a total size of 2 TByte are stored for 5 days on the network-attached disks of the HLRB II (R_4) in form of network attached storage (NAS) for short-term access, whereas 5 TByte of infrequently used simulation data are archived for 360 days by means of a storage area network (SAN) (R_5).

Further functionality of the virtual simulation service VS_1 offered to the customer includes graphical representation of simulation results by means of a visualization service. Due to the fact that the user has specific requirements regarding simulation data visualization, a customized visualization service (VS_4) provisioned by an external provider (VO_2) is used in order to visualize the simulation results by using the real services S_3 and S_4 which are each based on specialized visualization hardware or software (R_6 and R_7) offered at an external Grid service provider. In order to perform a rendering of three-dimensional turbulent flow graphics, the visualization service VS_4 is utilized for the time period of 2 hours. The accordingly resulting total costs are not directly obtainable by VO_1 since VO_1 does not have access to detailed accounting and charging records of VO_2. Instead, aggregated and consolidated pricing information is forwarded to VO_1 in form of a bill.

3.2 Accounting Model Application Methodology

Applying an extensive and flexible accounting model to a complex environment requires an elaborate methodology to be in place. Figure 7 provides an overview of the chosen model application methodology. It is structured into two main, chronologically separated building blocks, namely ABC taking input values from TCAS (0) and IT product cost calculation (1). IT product cost calculation relies on those activity costs determined by ABC. Section 3.2.1 and Section 3.2.2 explain procedures required for (0), while Section 3.2.3 details (1).

3.2.1 Annual Cost Input from TCAS

ABC seeks to identify costs per activity. In the applied methodology, activities are grouped by the criterion whether they can be related to an IT product (2) or they lack a product relation (3). Activities with product relation are further grouped in production activities (4) and activities that support production (5). The first category covers activities as determined by resource-specific instantiations of the introduced service constituent parts, namely *Processing*, *Storage*, *Transferring*, *Output*, *External*, and *Other*. The latter includes activities such as IT service and infrastructure management. Activities without product relation typically embrace facility management and administrative tasks (6).

Table 2 Considered Resource Attribution Keys

Attribution Key	Unit
Floor space consumed by a resource, including space required for maintenance	m^2
Annual resource power consumption	kW/year
Annual resource uptime	h/year

The accounting model takes annual costs of various types as input. These cost elements constitute typical values of TCAS. In the area of production-oriented activities, input values are needed in terms of annual costs with infrastructure performance (A). This is due to the fact that IT production in this context means the provisioning and composition of electronic services, such as a storage service. These services, out of which the final IT product is composed, are provided on infrastructure, that is, on IT resources. A given annual cost element with infrastructure performance is either attributed directly to the specific resource it relates to (I) or—in case these costs are not directly attributable to one of the existing IT resources—that cost element needs to be attributed indirectly by means of an allocation base, which is bound to an additional cost-relevant characteristic (II). IT resources, thus, reflect a concept from TCAS, namely the idea of a cost center. These cost centers embrace LRZ-internal computing and storage resources (C) as described in full detail in Section 3.1.1.

In order to allocate indirect costs to resources, attribution keys need to be in place as an allocation base. Table 2 lists those three attribution keys considered, namely floor space, power consumption, and uptime. The Grid accounting model is by no means limited to this specific set of attribution keys. This selection reflects information available at the LRZ, cost-wise relevant to the specific LRZ resources. The initial investment (in €, not differentiating between state and LRZ financing share) and annual operation costs (in €/year) for air conditioning infrastructure, emergency system, network infrastructure, and buildings constitute those LRZ cost elements with infrastructure performance that are not directly attributable to one of the considered computing or storage resources. As internal and external network traffic specific to Grid services is currently not separable from other traffic at the LRZ, all network-related costs need to be handled as indirect costs, even though, in principle, these costs would qualify to be directly attributable to network resources and, in a second step, to the according *Transferring* service constituent parts.

Table 3 lists directly attributable costs with infrastructure performance. These consider the annual cost elements available from LRZ's TCAS. Annual investment shares are not directly available, but calculated as the division of an

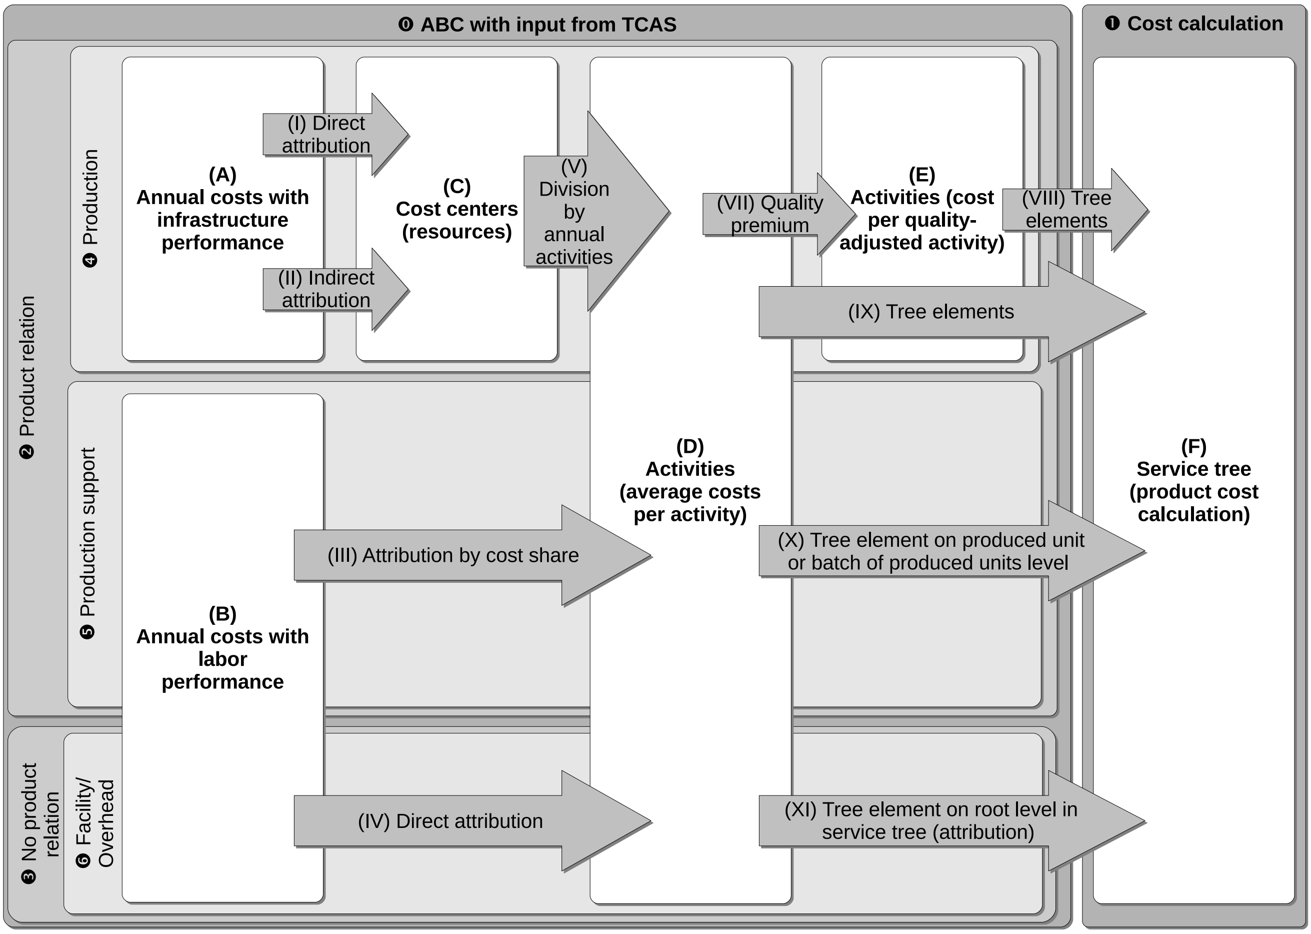

Fig. 7 Accounting Model Application Methodology Overview

Table 3 Directly Attributable Annual Costs with Infrastructure Performance

Cost Element	Unit
Annual investment share (reflects annual depreciation depending on initial resource investment and resource life time)	€/year
Annual electricity consumption (depending on the kWh price for electricity; excluding air conditioning)	€/year
Annual electricity consumption for air conditioning (depending on the kWh price for air conditioning)	€/year
Annual resource rental fee (applicable if resource is rented)	€/year
Annual software rental fee (total amount of software rental fees attributable to a resource)	€/year
Annual external labor costs (*e.g.*, for on-site service)	€/year
Annual material costs	€/year

IT resource's initial investment by its life time. Similarly, costs for annual electricity are calculated with the help of additional parameters. They result from multiplying an IT resource's annual uptime by its applicable kWh price and nominal power consumption.

After direct (I) or indirect (II) attribution of annual costs with infrastructure performance (A), total annual costs per considered IT resource—each representing a cost center—are revealed (C). Total annual costs per resource are defined as the sum of all direct annual cost elements and all indirect annual cost elements. The latter is attributed according to the respective annual cost share for air conditioning, emergency system, network infrastructure, and building costs. For instance, the annual air conditioning cost share for the Opteron cluster resource (*cf.* Section 3.1.1) is calculated by adding annual air conditioning operation costs to the annual air conditioning depreciation share (*i.e.*, the division of the original investment in air conditioning infrastructure by its life time), and multiplying this sum by the ratio of the Opteron cluster's nominal power consumption to the total nominal power consumption of all considered resources.

In contrast to annual costs with infrastructure performance (A), annual costs with labor performance (B) do not require an intermediate attribution step to cost centers, *i.e.* resources, since labor performance costs are directly related to activities (D). Annual costs with labor performance (B) and production support (5) embrace human labor activities which are grouped after process activities of the IT Infrastructure Library (ITIL) [28] version 2. These best practices determine the de-facto standard in service management. The respective books on infrastructure and service managements are of particular importance for this work as they are concerned with production support activities. Due to the fact that the LRZ cannot provide any information on employee work assignments for legal reasons, an estimation of which ITIL activity is more costly than another is not feasible at this time. Therefore, it is assumed initially that all ITIL activities need to cover an equal cost share. These relative cost shares (20% for each ITIL activity, since 5 ITIL activity types are considered) are used as keys to attribute (III) annual costs with labor performance (B) and production support (5) to the respective ABC activities (D). Annual costs are available at the LRZ for two labor categories, internal operations and internal support. For both categories, the number of positions at the LRZ is multiplied by the average wage, the results added, then multiplied by the applicable percental cost share, and finally divided by the mathematical product of annual working days and daily working hours. By this calculation (III), average costs per hour are gained for each considered ITIL activity (D).

Annual costs with labor performance (B) without product relation (3) include facility management and administrative overhead activities. For both types, average costs per activity (D) are directly retrievable (IV), *i.e.*, an attribution according to a key is not necessary. Consequently, the applied calculation method represents a simplified version of the method used for ITIL activities: The number of positions at the LRZ is multiplied by the average wage, and the result is divided by the mathematical product of annual working days and daily working hours. This results in average costs per hour and activity (D), whereas these activities embrace the mentioned facility management and administrative overhead.

3.2.2 Resource-specific Activity-based Costing

Table 4 gives an overview of those 15 activities (D) resulting from either dividing resource-attributed costs by annual activities (V) or attributing annual costs with labor performance (B) by either cost share (III) or by direct attribution (IV). For each activity, the corresponding service constituent part is listed. Production activities (4) are represented by a *Processing*, *Storage*, or external *Output* service constituent part, while production support (5) and facility/overhead activities (6) are represented by the service constituent part *Other*.

This list of activities constitutes the key functional step in applying the Grid accounting model as it comprises those activities that form the basis for ABC. At this step in model application (D), first the full list of activities for building a service tree (F) from is available, and second the average costs per activity are revealed. This means, *e.g.*, for the *Processing* activity Altix that costs for computing on that resource per CPU second are known. In general, costs per activity and the accordingly applicable metric are determined.

Table 4 Activities and Service Constituent Parts

Activity	Service Constituent Part
HLRB II	Processing
32 Bit	Processing
IA 64	Processing
Opteron	Processing
Altix	Processing
Backup, archive, SAN	Storage
NAS	Storage
VR cluster	Output (external)
RV cluster	Output (external)
IT infrastructure design and planning	Other
IT infrastructure deployment	Other
IT infrastructure operations	Other
IT infrastructure technical support	Other
Facility management	Other
Administrative overhead	Other

All *Processing* and *Output* activities use CPU seconds, all *Storage* activities use resource reservation events, and all *Other* activities use working hours as metric. As *Output* activities are not provided internally, but are offered by an external provider (see Section 3.1.2 for scenario details), cost calculation and metric selection decisions lie within that other organization's responsibility. Calculations for these activities, hence, are not performed with the same granularity as it is the case for internal activities. Consequently, the metric of CPU seconds is not used for actual cost calculations, but seen as a metric to appear on a bill received by that other organization.

From a business logic viewpoint, metrics are bound to ABC's activity drivers. Activity drivers are perceived as the event or fact that influences an activity's intensity with respect to costs incurred. For *Processing* activities, this cost triggering event is found, for instance, in the atomic computing activity of a CPU second used on a given resource. Those chosen metrics, however, are neither fully deterministically selected nor are they elements of a statically defined set of available metrics. Accordingly, those metrics chosen here are on the one hand inspired by the overview on accountable units provided in Figure 5, on the other hand determined by metering capabilities available at the LRZ.

The activities determined as shown in Table 4 can either directly (VIII to XI) form elements of the service tree (F) for product cost calculation (1) or, before that, they can be further refined in order to support quality adjustments (E). Quality-adjusted activities are determined for all internal activities, thus according to the applicable scenario (*cf.* Section 3.1.2), for all *Processing* and *Storage* activities. The underlying principle for quality adjustments funds on a quality premium scheme. It supposes that non-adjusted activities (D) include a standard configuration. For *Storage* activities, a two-dimensional standard configuration is assumed. For backup, this includes a resource reservation of 1 TByte capacity for the duration of 360 days, while for NAS, a capacity of 1 GByte for the duration of 30 days is assumed. Similarly, *Processing* activities see a presumed two-dimensional standard configuration of 1'024 CPUs with 4 GByte of main memory per CPU available in case of HLRB II, and of 32 CPUs with 1 GByte main memory per CPU for all other LRZ computing resources. Whenever a standard configuration needs to be changed increased costs for (potentially) intensified resource usage are possible to be reflected by a cost premium (VII), which is a percental supplement to the average activity costs (D).

A quality premium is represented by ABC's resource driver concept. Resource drivers—as opposed to activity drivers—are events or facts that influence a resource's usage intensity, such as a resource reservation for extended storage capacity. Multi-dimensional quality premiums are implemented by defining a multi-dimensional unit. For *Storage* activities, that is GBd (GByte day), while for *Processing* activities, a unit called CGB (CPU second GByte) is used. Both units are calculated as the mathematical product of each involved single-dimension unit. For instance, quality-adjusted costs for the *Storage* activity NAS are determined by dividing first the standard, *i.e.*, not quality-adjusted cost for NAS by its standard GBd configuration (368'640 GBd as the mathematical product of 1'024 GByte and 360 days). This intermediate result is multiplied by the respective quality premium, resulting in quality-adjusted costs measured by a unit of €/GBd.

While the same quality premium concept applies for calculation of quality-adjusted activity costs of *Storage* activities and of *Processing* activities, the respective used multi-dimensional units need to be differentiated clearly: The unit of GBd is used exclusively for *Storage* activities, and CGB is used exclusively for *Processing* activities.

3.2.3 IT Product Cost Calculation

According to the scenario-specific service tree depicted in Figure 6, in the following the methodology introduced in Section 3.2 concerning product cost calculations is altered by means of concrete values. On the one hand used data directly correlates to some extent to concrete values and parameter settings acquired from the LRZ, on the other hand some of the data is based on assumptions or approximations, respectively.

On top-level, the virtual simulation service VS_1 offered within VO_1 is composed of virtual services being represented by the service constituent parts *Processing* (VS_2), *Storage* (VS_3), as well as the service constituent part *Output* (external), reflecting the usage of the virtual service VS_4 offered at an external Grid service provider. Additionally, in order to adequately reflect the activities being performed using the virtual simulation service, tasks with regard to the design and planning of the compound virtual service VS_1 have to be taken into consideration as well, resulting in a total of 10 working hours estimated which are being mapped on a batch of 20 service requests. This implies that 5 percent of the resulting costs for these activities have to be calculated per service invocation. Additionally, costs occurring with respect to facility management (0.5 hours per service request assumed) as well as administrative overhead (1 hour per service request estimated) being covered by the service constituent part *Other* also have to be incorporated as relevant activities having a direct relation to the compound virtual simulation service VS_1. Finally, expenses originating from activities with respect to IT service management have to be taken into consideration as well. Due to the high degree of dynamics within the context of DVOs as well as rapidly changing business processes, concerning the compound virtual simulation service VS_1, configuration management and change management constitute important ITIL activities which result in 15 working hours estimated each, also being mapped on a batch of 20 service requests. These subcategories of IT service management, thus, result in total in 30 working hours per 20 service requests.

The virtual computation service VS_2 itself is performed using the composed virtual computation resource VR_1 comprising the HLRB II (R_1), the IA 64 cluster (R_2) as well as the Altix cluster (R_3). Within the scenario, 512 processors of the HLRB II are used for 2.5 hours (=9'000 CPU seconds) each with an average main memory utilization of 2 GByte per processor, resulting in 1'024 CGB which is lower than the standard configuration of 512 CPUs and 4 GByte of reserved main memory by the factor of 2. Additionally, in order to process the user job, the entire IA 64 cluster (R_2) comprising a total of 220 processors is utilized for 4 hours (=14'400 CPU seconds) along with an average memory usage of 1 GByte per CPU resulting in 220 CGB in total. Finally, 25 percent (=32 processors) of the Altix cluster are utilized for a time period of 6 hours (=32'600 CPU seconds) each, together with the utilization of 1.5 GByte of main memory per CPU (=48 CGB) which exceeds the standard configuration for computing resources, thus, resulting in quality-adjusted costs per activity. Moreover, concerning the virtual computation resource VR_1 costs regarding the IT infrastructure deployment as well as the IT infrastructure operations have to be taken into consideration, resulting in a total of 10 working hours estimated per activity and per month which have to be mapped on a batch of 5 service requests of the virtual computation service VS_2. Due to fact that negotiated QoS parameters with respect to execution time have to be met (*cf.* Section 3.1.2), also costs reflecting ITIL activities in relation to Service Level Management (SLM) resulting in 1 working hour estimated per service request have to be incorporated into the product cost calculation.

Additionally, in the scenario the compound virtual storage service VS_3 provided by VO_1 comprises two real data services S_1 and S_2 offered at the LRZ which in turn are provisioned upon the physical storage resources R_4 in form of a network attached storage and R_5 being a storage area network. Within the scenario depicted in Section 3.1.2, the real data service S_1 is used in order to store frequently used simulation results with a size of 2 TByte for the time period of 5 days, which results in a total of 10'240 GBd, thus exceeding the standard capacity and duration activity for storage resources. Besides, in order to archive 5 TByte of simulation data on the long-term data storage for 360 days, the real data service S_2 making use of a magnetic tape system (R_5) offered at the LRZ is used. The utilization of the real long-term storage service S_2 results in a total of 1'843'200 GBd. In addition, costs reflecting the IT infrastructure technical support of the storage resources have to be considered when calculating the costs of the activities being performed by means of the virtual data service. Hence, overall costs of 5 working hours estimated in relation to technical storage resource support—to be mapped on a batch of 10 service invocations—also have to be calculated per service request. In order to assure long-term archival storage of the simulation data using the virtual storage service VS_3, activities with respect to continuity management also have to be considered, resulting in 0.5 working hours estimated per service request.

Finally, as shown in the service tree presented in Section 3.1.2, a virtual visualization service (VS_3) is part of an external Grid service provider (VO_2) and is used in order to graphically represent obtained simulation results by consuming two real visualization services, S_3 and S_4. According to the bill which is forwarded by the external Grid service provider to the customer VO_1, a VR cluster (R_6) as well as a remote RV cluster (R_7) both represented by the service constituent part *Output* (external) are each utilized for 1 hour (=3'600 CPU seconds).

3.3 Key Evaluation Objectives and Requirements

Based on the fact that the identified activities are resource-specific and have to be adapted to the particular resources they reflect (*cf.* Section 2.3), the evaluation of the proposed accounting model needs to include a detailed infrastructure and service analysis. This analysis needs to document what resources are available (formally also reflected by resource-specific activities) and what commercial services need to be run on them (leading to a bill of activities and the fully documented service tree). Based on this information, the evaluation shall reveal what costs need to be covered per service request.

As input data to the Grid accounting model, information from the traditional financial and the cost accounting—both areas of economic (as opposed to technical) accounting—is needed. This comprises, for instance, information on investments or maintenance costs incurred during a fiscal year. These cost elements are first categorized into cost categories and secondly either directly or indirectly allocated to cost centers. Those steps still determine typical activities in a traditional accounting system. The evaluation, thus, needs to answer the questions whether such information was available at the LRZ in the first place and if it was of the right granularity in order to deliver meaningful input for the accounting model.

Overall, the conducted evaluation shall answer how well the existing Grid accounting model is able to calculate costs to be covered for a specific service request. In particular, and by means of varying assumptions, the evaluation shall depict for a real Grid infrastructure what input data and also what level of detail is required to allow the model to produce meaningful results with reasonable costs incurred by using the model. Further, potential improvements to the model need to be derived. Driven by these key evaluation requirements outlined, the set of specific qualitative evaluation criteria is determined as listed subsequently:

- **Model functionality**: General functionality of the Grid accounting model and information content provided is assessed. This comprises in particular the achieved level of result expressiveness, addressing both, gained insight as well as limitations encountered.
- **Model parametrization**: The applied set of service constituent parts, considered metrics, and chosen activity/resource drivers is examined in detail. This addresses unit characteristics with associated interdependencies. Effects of changes in calculation input parameter assumptions are of particular interest.
- **Model application context**: The respective available input data for model application by means of the presented multi-provider scenario is assessed. Sensitivity analyses with respect to product cost impact caused by scenario parameter changes are evaluated.

The discussion on model functionality is conducted in Section 5.1, while Section 5.2 assesses results with respect to model parametrization, and Section 5.3 is concerned with an evaluation of the model application context.

4 Results

Driven by the outlined application and evaluation methodology, the proposed Grid accounting model for DVOs is ap-

Annual Costs

Infrastructure Performance

		Air conditioning	Emergency system	Network infrastructure	Building	Total	Unit
Indirect attribution	Investment	0	0	0	46'000'000	46'000'000	€
	Annual operations	144'836	100'735	2'200'000	242'120	2'687'691	€/year
Additional cost-relevant characteristics	Life time	6	6	3	25		years
	Attribution key	Power consumption	Floor space incl. maintenance	Floor space incl. maintenance	Floor space incl. maintenance		

Cost centers (Resources)

		HLRB II	32 Bit	IA 64	Opteron	Altix	Backup, archive, SAN	NAS	Total	Unit
Attribution keys	Floor space incl. maintenance	258	16.25	16.25	16.25	16.25	170	10	503	m2
	Power consumption	1100	30	50	40	25	85	30	1'360	kW
	Uptime	8'590	8'670	8'670	8'670	8'590	8'760	8'760	60'710	h/year
Additional cost-relevant characteristics	kWh price (electricity)	0.11	0.11	0.11	0.11	0.11	0.11	0.11		€/kWh
	kWh price (air conditioning)	0.044	0.044	0.044	0.044	0.044	0.044	0.044		€/kWh
	Life time	5	3	5	3	5	5	4		years
	CPUs	9'728	134	220	192	128	n/a	n/a		
	Resource investment	54'000'000	229'355	555'060	190'000	1'200'000	6'300'000	500'000	62'974'415	€
Direct attribution	Annual investment share	10'800'000	76'452	111'012	63'333	240'000	1'260'000	125'000		€/year
	Electricity	1'039'390	28'611	47'685	38'148	23'623	81'906	28'908	1'288'271	€/year
	Air conditioning	415'756	11'444	19'074	15'259	9'449	32'762	11'563	515'308	€/year
	Resource rental	0	0	0	0	0	0	0	0	€/year
	Software rental	0	0	0	0	0	0	0	0	€/year
	External labor	0	0	0	0	109'480	0	0	109'480	€/year
	Material costs	0	0	10'000	0	0	0	0	10'000	€/year
Annual costs	Direct annual costs	12'255'146	345'862	742'831	306'741	1'582'552	7'674'668	665'471		€/year
	Annual air conditioning cost share	117'147	3'195	5'325	4'260	2'662	9'052	3'195		€/year
	Annual emergency system cost share	51'669	3'254	3'254	3'254	3'254	34'046	2'003		€/year
	Annual network infrastructure cost share	1'128'429	71'074	71'074	71'074	71'074	743'539	43'738		€/year
	Annual building cost share	1'067'966	67'265	67'265	67'265	67'265	703'699	41'394		€/year
	Total annual costs	14'620'358	490'650	889'749	452'594	1'726'807	9'165'004	755'800	28'100'962	€/year

Labor Performance

		Internal operations	Internal support	Unit
Indirect attribution	Wage	82'574	82'574	€/year
	Positions	5.5	10	
Additional cost-relevant characteristics	Attribution key	Cost share	Cost share	
	Annual working days		220	days/year
	Daily working hours		8	hours/day

		IT infrastructure design and planning	IT infrastructure deployment	IT infrastructure operations	IT infrastructure technical support	IT service management	Unit
Attribution keys	Cost share	20	20	20	20	20	%

		Internal facility management labor	Internal administration labor	Unit
Direct attribution	Wage	0	52'534	€/year
	Positions	0	2	
Additional cost-relevant characteristics	Annual working days		220	days/year
	Daily working hours		8	hours/day

Fig. 8 Annual Costs Calculation

plied to the determined multi-domain Grid accounting scenario. This is achieved by a full-cost calculation performed with input data from the LRZ.

Figure 8 presents annual cost calculations which include indirect costs resulting from the LRZ air conditioning system, the emergency system, its network infrastructure as well as building costs. It needs to be stressed that initial investments in the first three mentioned categories are subsumed in the initial investment amount of the LRZ building. Thus, a zero investment value for, *e.g.*, the emergency system reflects the fact that these investment costs are not separately obtainable.

While those investment and annual operations infrastructure costs reflect indirect costs (II in Figure 7), Figure 8 also depicts direct costs (I in Figure 7) such as material costs where applicable. Direct and indirect annual costs are attributed to the respective set of LRZ IT resources, consisting of computing infrastructure like the HLRB II cluster and of storage infrastructure such as NAS. These LRZ resources serve as cost centers (C in Figure 7) that need to bear annual costs of approximately 28 million €.

Furthermore, Figure 8 visualizes annual costs with labor performance (B in Figure 7). This covers in particular LRZ-specific information on number of positions, wages, working days, and working hours. It needs to be stressed, however, that these numbers are simplified target figures so that, in reality, differing numbers might apply. Additionally and similar to those zero investments reported for, *e.g.*, the LRZ air conditioning system, figures for internal facility management labor are zero. This is due to the fact that facility management costs are included in the respective number for annual building operations. Annual facility management labor costs—although being reported as zero here—and annual administration labor costs are directly attributed (IV in Figure 7) to activities, whereas annual operations and support costs are assigned (III in Figure 7) to activities by means of an (equal) cost share of 20%.

Figure 9 focuses on activity-related cost calculations (D in Figure 7) of both considered activity cost types, average costs per activity and—with regard to non-standard activity configurations—quality-adjusted activity costs (E in Figure 7). The calculation of average activity costs for activities of type *Processing* bases on the assumption that all LRZ computing resources show a capacity utilization of 80%. For the time being, the exact capacity utilization value is not measured at the LRZ so that it needs to be estimated. A value of 80% determines a conservative estimation, since annual usage statistics at the LRZ show long queues of waiting jobs. These statistics are considered for all computing resources other than the HLRB II cluster. This cluster has seen a major increase of nodes in 2007 from 4'096 to 9'728 CPUs—a fact which does not become apparent in the annual usage figures. In addition, annual statistics only account for the aggregated uptime of so-called batch nodes (a logic composite of currently 512 CPUs). Thus, annual statistics for the HLRB II cluster do not allow to estimate its capacity utilization level reliably. For that reason, the same level of 80% is assumed for HLRB II activities.

Average costs for *Storage* and *Output* activities in Figure 9 determine estimated values. In the case of *Storage*, these values are estimated from previous LRZ experience. *Output* activities for visualization of results represent external activities which are provided by VO_2 (*cf.* Section 3.1.2). The according activity costs constitute costs from the viewpoint of VO_1 only, whereas from VO_2's viewpoint, they constitute billed values. Billing information might not only cover VO_2's production costs, *i.e.*, it might not follow a strict cost-oriented pricing, but incorporate a pricing scheme which is profit maximizing. In addition, visualization services are run on highly specialized, expensive equipment. For these reasons, average *Output* activities costs are estimated to be higher than, *e.g.*, internal computing activity costs.

Standard duration and capacity for *Storage* activities as well as standard CPU and main memory numbers determine estimated values from LRZ experience, adopted to the presented Grid scenario. The according quality premium values cannot be substantiated at this time by specific statistics on resource drivers and, thus, costs caused by providing non-standard resource configurations. Therefore, quality premiums are initially set to an assumed (low) percentage of 5%.

Figure 10 visualizes product cost calculations according to the service tree depicted in Figure 6. These calculations multiply the respective activity costs as outlined in Figure 9 by the applicable accounted or billed units as described in Section 3.1.2 (scenario definition) and Section 3.2.3 (product calculation specifics). This results in monetary values representing costs incurred by each activity and, in sum, in total product costs of 4'656 €. The virtual service VS_1 in Figure 6 relates in this context to the product for which costs are calculated. Thus, in application of the outlined methodology of an activity-based, resource specific, full cost-oriented Grid accounting model, this calculation determines those costs that need to be covered by each invocation of VS_1. It needs to be stressed that the resulting amount reflects costs, which are not to be mistaken for product pricing.

5 Discussion

Based on the evaluation objectives outlined in Section 3.3, this section assesses the results gained from the Grid accounting model application by means of the presented cost calculation. This implies the results discussion regarding model functionality (*cf.* Section 5.1), possibilities of model parametrization (*cf.* Section 5.2), and the according evaluation of the model application context (*cf.* Section 5.3).

Activities

Product Relation — Production — Average costs

	Processing	Storage	Transferring (internal)	Transferring (external)	Output (external)	Other	Unit
Activity driver	Computing	Resource reservation	TCP/UDP traffic	TCP/UDP traffic	Computing	Labor	
Metric	CPU seconds	Resource reservation events	TCP/UDP segments	TCP/UDP segments	CPU seconds	Working hours	

	HLRB II	32 Bit	IA 64	Opteron	Altix	Backup, archive, SAN	NAS	VR cluster	RV cluster	Unit
Service constituent part	Processing	Processing	Processing	Processing	Processing	Storage	Storage	Output (external)	Output (external)	
Activity driver	Computing	Computing	Computing	Computing	Computing	Resource reservation	Resource reservation	Computing	Computing	
Metric	CPU seconds	CPU seconds	CPU seconds	CPU seconds	CPU seconds	Resource reservation events	Resource reservation events	CPU seconds	CPU seconds	
Effective annual computing activities	240'662'937'600	3'345'926'400	5'493'312'000	4'794'163'200	3'166'617'600	n/a	n/a	n/a (external)	n/a (external)	CPU seconds
Effective annual storage activities	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a (external)	n/a (external)	Resource reservation events
Average costs per activity	0.00006075	0.00014664	0.00016197	0.00009441	0.00054532	250	1	0.03000000	0.02000000	€/metric

Storage quality adjustments

	Backup, archive, SAN	NAS	Unit
Duration and capacity resource driver	Altered duration and/or capacity reservation	Altered duration and/or capacity reservation	GBd (GByte day)
Standard duration activity	360	30	Day
Standard capacity activity	1'024	1	GByte
Altered duration and/or capacity quality premium	5	5	%
Quality-adjusted costs per activity with altered duration and/or capacity	0.00071208	0.03500000	€/GBd

Processing quality adjustments

	HLRB II	32 Bit	IA 64	Opteron	Altix	Unit
CPU and main memory resource driver	Altered CPU and/or main memory reservation	Altered CPU and/or main memory reservation	Altered CPU and/or main memory reservation	Altered CPU and/or main memory reservation	Altered CPU and/or main memory reservation	CGB (CPU second Gbyte)
Standard CPU activity	512	32	32	32	32	CPU
Standard main memory activity (per CPU)	4	1	1	1	1	GByte
Altered CPU and/or main memory premium	5	5	5	5	5	%
Quality-adjusted costs per activity with altered CPU and/or main memory	0.00001595	0.00015397	0.00017007	0.00009913	0.00057258	€/CGB

Production support — Average costs

	IT infrastructure design and planning	IT infrastructure deployment	IT infrastructure operations	IT infrastructure technical support	IT service management	Unit
Service constituent part	Other	Other	Other	Other	Other	
Activity driver	Labor	Labor	Labor	Labor	Labor	
Metric	Working hours	Working hours	Working hours	Working hours	Working hours	
Average costs per activity	145.44	145.44	145.44	145.44	145.44	€/hour

	Facility management	Administrative overhead	Unit
Service constituent part	Other	Other	
Activity driver	Labor	Labor	
Average costs per activity	0.00	59.70	€/hour

Fig. 9 Activities Calculation

5.1 Model Functionality

Both, the methodology developed and the specific calculation performed reflect a high level of expressiveness. This is particularly substantiated by a most direct implementation of the set of key Grid accounting model characteristics: The calculation incorporates annual costs resulting from the relevant LRZ infrastructure and IT resources, which act as a cost center from an (economic) accounting viewpoint. This principle of resource-specific calculations is continued by the definition of resource-adapted activities. These activities are not only resource-specific, but support another important model characteristic as they are quality-aware. The introduced quality premium approach allows for configuring

IT Product Cost Calculation

Simulation

		Accounted CPU seconds per CPU	*CPUs*	*Main memory per CPU*	*Activity costs*	*Unit*
Processing	*HLRB II*	9'000	512	2	146.97	€
	IA 64	14'400	220	1	538.78	€
	Altix	32'600	32	1.5	895.98	€
		Duration	*Capacity*	*Activity costs*		
Storage	*NAS*	5	2'048	358		€
	Backup, archive, SAN	360	5'120	1'313		€
		Billed CPU seconds	*Activity costs*			
Output (ex-ternal)	*VR cluster*	3'600	108.00			€
	RV cluster	3'600	72.00			€
		Accounted working hours	*Unit/batch activity map-ping factor*	*Activity costs*		
Other	*IT infrastruc-ture design and planning*	10	0.05	73		€
	IT infrastruc-ture deploy-ment	10	0.20	291		€
	IT infrastruc-ture opera-tions	10	0.20	291		€
	IT infrastruc-ture technical support	5	0.10	73		€
	IT service management (standard)	30	0.05	218		€
	IT service management (quality-ad-justments)	1	1.00	145		€
	IT service management (continuity management)	0	1.00	73		€
		Accounted working hours	*Activity costs*			
Other	*Facility man-agement*	0.5	0.00			€
	Administrat-ive overhead	1	59.70			€
Product costs	4'655.86					€

Fig. 10 Product Cost Calculation

non-standard offers according to user demand—while still being able to express increased resource usage or even losses incurred by resources that might not be attributable to other users even though they are not used by the initial user. For instance, main memory for a node of the HLRB II cluster may be limited for one user to 2 GByte. The remaining 2 GByte (4 GByte is standard per node), however, will not be available for another user. In that light, the existence of a quality premium seems appropriate.

In a similar way, the calculation has proven the Grid accounting model's theoretical nature of being highly parameterizable and, thus, being flexible, extensible, and generically applicable. Flexibility is reflected exemplary by a high degree of freedom to define input parameters, such as attribution keys. Extensibility is visualized by the example of freely configurable standard activities and quality adjustments. General applicability is substantiated exemplary by the fact that costs related to *Transferring*—as one of those four basic service constituent parts of the original Grid accounting model—could be handled as an element of TCAS for pragmatic LRZ-specific reasons, even though *Transferring* activities were foreseen initially to constitute a central element in the ABC part of the calculation.

The developed methodology and the appropriately determined calculation are found to first integrate successfully the respective viewpoints of technical and economic accounting. Secondly, they show that the Grid accounting model's expressiveness finds implementation in a practically viable way to determine product costs for multi-domain Grid service scenarios. For the considered scenario, product costs of 4'656 € were calculated, out of which a share of 34% resulted from *Processing* costs, a comparably high share of 36% from *Storage* costs, a 4% share from *Output* costs, and a 26% share from *Other* costs. At first glance, costs of 4'656 € per service instantiation might seem to be relatively high. However, the cost/performance ratio has to be considered in relation to the respective field of application (*e.g.*, consider an automotive manufacturer within a fully commercial environment).

Although the calculation demonstrates a successful Grid accounting model applicability in general, it sees potential for further improvements. For instance, it does not consider load balancing aspects which might be of high impact for a supercomputing environment. Similarly, the calculation as it stands needs to consider costs caused by unused but not attributable resources in a more fine-granular way. This means that the concept of quality premiums needs to be extended in order to better support competition for resources.

Furthermore, the calculation has revealed that the proposed Grid accounting model is in its application to a real-world environment like the LRZ not fully transparent for a model user. In-depth knowledge, both about the model itself as well as the underlying infrastructure and service parameters is still needed. Thus, model and calculation should be extended to define, *e.g.*, the generally applicable, relevant set of technical accounting metering points.

In order to conclude, the calculation is found to provide valuable results in product cost determination by implementing the generic Grid accounting model in its full expressiveness and successfully applying it to a real-world environment. However, model application requires at this time considerable effort in configuring and parametrizing the calculation.

5.2 Model Parametrization

As the Grid accounting model was applied to a real-world environment for the first time, a number of calculation parameters were required to be estimated. Other parameters, such as those mentioned as *Transferring* costs, could not be metered in a way that would have allowed for data usage as initially intended by the model. Despite such practical concerns, the resulting calculation is found to constitute an extensive and effective model application case. In the case that assumptions were taken, these could be either estimated

Calculation Input Parameter Sensitivity Analysis

	Building	Air conditioning	Emergency system	Network infrastructure
Investment	0.2	<0.1	<0.1	0.2
Annual operations	<0.1			
Life time	0.1			

	HLRB II	IA 64	Altix	Backup, archive, SAN	NAS	VR cluster	RV cluster
Floor space	0.1	0.2	0.1	0.1	<0.1		
Power consumption	<0.1	<0.1	<0.1	<0.1	<0.1		
Uptime	0.3	1	1.7	<0.1	<0.1		
kWh price (electricity)	<0.1	<0.1	<0.1	<0.1	<0.1		
kWh price (air conditioning)	<0.1	<0.1	<0.1	<0.1	<0.1		
Life time	0.2	0.1	0.2	<0.1	<0.1		
CPUs	0.3	1.1	1.7				
Resource investment	0.2	0.9	1.6	<0.1	<0.1		
Resource utilization	0.3	1.1	1.7				
Average costs per activity				2.8	0.8	0.2	0.2
Standard main memory activity	0.3	1.1	1.7				
Standard duration activity				2.6	0.7		
Standard capacity activity				2.6	0.7		
Altered duration and/or capacity quality premium	<0.1	<0.1	0.1	0.1	<0.1		

	Internal operations	Internal support	Internal administration labor
Wage	0.9	1.6	0.1
Positions	0.9	1.6	0.1

	IT infrastructure design and planning	IT infrastructure deployment	IT infrastructure operations	IT infrastructure technical support	IT service management
Cost share	0.4	0.2	0.2	0.4	0.5

Fig. 11 Percental Impact on Product Costs of a 10% Calculation Input Parameter Change

from past LRZ experience or they were clearly termed as assumptions.

In the light that some calculation parameters were estimated or assumed, a sensitivity analysis of key parameter changes helps to assess one parameter's change impact to the overall product cost calculation. Figure 11 documents the respective percental change in product costs of (initially) 4'656 € if one calculation input parameter is changed by 10% of its value, *ceteris paribus*, meaning that all other parameters are left unchanged. Most caused changes are assessed marginal with an impact on product costs of less than 0.1%. However, there is a considerable impact on product costs in some areas. The top five impact areas are identified as follows: Changes of 10% on Backup, archive, SAN parameters of either average costs per activity, standard duration activity, or standard capacity activity result in a change in product costs in the range of 2.6-2.8%. In other words, these parameter changes are leveraged by about a fourth. The second most important product cost change (in the range of 1.5-1.7%) is observed when selected parameter values for the Altix cluster are altered by 10%. Changes on wage or position values for internal support labor fall into a comparable class of relative impact, namely of 1.6%. The third largest leverage effect show selected parameter changes for the IA 64 cluster, closely followed by effects incurred by parameter changes in the area of internal operations labor. The fifth largest impact on product costs show parameter changes for the NAS infrastructure, such as for average costs per activity (in the range of 0.7-0.8%).

Scenario Parameter Sensitivity Analysis

	HLRB II	IA 64	Altix	Backup, archive, SAN	NAS	VR cluster	RV cluster
Accounted CPU seconds per CPU	0.3	1.2	1.9				
CPUs	0.3	1.2	1.9				
Main memory per CPU	0.3	1.2	1.9				
Duration				2.8	0.8		
Capacity				2.8	0.8		
Billed CPU seconds						0.2	0.2

	IT infrastructure design and planning	IT infrastructure deployment	IT infrastructure operations	IT infrastructure technical support	IT service management	Administrative overhead
Accounted working hours	0.2	0.6	0.6	0.2	0.9	0.1
Unit/batch activity mapping factor	0.2	0.6	0.6	0.2	0.9	0.1

Fig. 12 Percental Impact on Product Costs of a 10% Scenario Parameter Change

It has to be stressed that these sensitivity analyses conducted cannot provide completely unbiased insight with respect to product cost impact due to inherent dependencies on the chosen scenario. For instance, any change on input parameters in relation to the 32 Bit cluster will not show any effect on product costs here, since this infrastructure is not considered to be used in the applicable scenario. Nevertheless, these sensitivity analyses allow to identify parameter values of particular importance which need careful inspection—especially in case such a parameter was assumed or estimated as it is the case, *e.g.*, for the average costs per activity for Backup, archive, SAN. Thus, this calculation cannot only be helpful for product cost calculations, but it can serve as an instrument for optimizations.

5.3 Model Application Context

The developed methodology and the resulting calculation both document that the Grid accounting model was successfully applied to existing LRZ infrastructure. The chosen scenario, however, incorporates specifics that do not reflect current LRZ characteristics. Most prominently, the LRZ does not offer at present a virtual service similar to VS_1. Neither are virtualized resources made available as Grid services in a multi-domain environment. For such reasons, the scenario chosen needs to be deemed to be of a partially artificial nature. In the same manner as those previously mentioned practical limitations of partially lacking technical accounting metering data, this bears a risk to lower overall calculation significance. Thus, a sensitivity analysis of scenario parameter value changes is of particular interest.

As Figure 12 depicts, these sensitivity analyses conducted for 10% scenario parameter value changes show on average a larger impact on product costs than the average percental impact caused by those calculation input parameter changes assessed in Figure 11. The respective top five impact areas are identified as follows: Changes in duration and capacity scenario parameters for Backup, archive, SAN cause the highest change in product costs (2.8%). This is followed by parameter changes to the Altix cluster (1.9%) and the IA 64 cluster (1.2%), respectively. IT service management parameter changes of 10% result in altered product costs of 0.9%, while duration and capacity parameter changes to NAS show an impact of 0.8% to product costs.

In accordance with those conclusions drawn in Section 5.1 and Section 5.2, these percental impact numbers consolidate the identified need for an improved, more fine-granular technical accounting that substantiates parameter values by means of metered data instead of assumed values.

In summary, these areas of future improvements with regard to the proposed Grid accounting model were identified in the course of the successful model application to the LRZ environment as performed and discussed so far:

- Consideration of load balancing aspects.
- Extension of the concept of quality premiums to better support competition for resources.
- Consideration of costs caused by unused but not attributable resources in a more fine-granular way.
- Definition and integration of generally applicable set of metering points for technical accounting.

6 Summary and Conclusions

With the ongoing trend of adopting Grid systems as a means for service-oriented computing in DVOs, the need for appropriate support mechanisms becomes apparent. Accounting of Grid resource and service usage determines the central support activity since it prepares accounting records that provide the main input for analysis, optimization, and in particular for charging and billing purposes.

An embracing study of existing Grid accounting systems revealed that these approaches focus primarily on technical precision and on project-specific issues, whereas they do not support multi-provider scenarios or virtualization concepts, nor are existing approaches based on appropriate economic accounting principles regarding cost calculation. Consequently, the determined resource-based, highly flexible accounting model for DVOs [15] combines both, technical and economic accounting by means of Activity-based Costing, service constituent parts and defined accountable units.

Driven by the successful preliminary conceptual evaluation of the proposed accounting model for DVOs, throughout this paper, a full-fledged evaluation of the presented approach has been undertaken. For this purpose, the generic accounting model was applied to an existing operational Grid infrastructure operated by the Leibniz Supercomputing Centre in Garching near Munich, Germany in order to reveal the key set of practical aspects relevant for this model's application and to determine potential model improvements and extensions respectively.

Therefore, based on a brief recapitulation of key mechanisms for Grid service accounting in DVOs, addressing the proposed DVO service and Grid accounting models, a taxonomy of Grid resources was developed, providing an appropriate basis for the identification of accounting units and metrics adequately reflecting resource consumption and service usage, hence, serving as valuable input with respect to the evaluation methodology.

In accordance with those identified accounting model characteristics, the appropriate methodology for the application and evaluation of the proposed model was specified in detail. This task included an in-depth investigation into the LRZ Grid infrastructure and provisioned Grid services as well as a description of financial, cost-related input data. Additionally, a multi-domain Grid accounting scenario, which was enhanced with concrete values and parameter settings, was introduced providing the basic principles for subsequent model application and evaluation tasks.

Based on the gained insights, various model calculations comprising an annual cost calculation, an activities calculation as well as a product cost calculation have been performed and discussed according to a set of previously identified evaluation criteria regarding model functionality, parametrization, and application context. In this regard, the assessment of those results gained from the presented cost calculation has revealed that the Grid accounting model constitutes an expressive, highly flexible, extensible as well as generically applicable tool for two inter-related key purposes, (a) Grid service cost calculation and (b) cost optimization identification.

The proposed Grid accounting model demonstrates its general applicability to various organizational contexts that may range from small and medium-sized enterprises to large supercomputing centers such as the LRZ. Due to the model's universal design putting emphasis on typical and configurable activities in a Grid environment, insights gained from the model application case at the LRZ are transferable to further environments. Those model application steps, *e.g.*, with respect to activity configurations, resource adaptations, and quality premium definitions performed, will be conducted methodologically fully in line with the application case performed. Thus, even though another organizational application context may expose different resources or other calculation input data from TCAS, the Grid accounting model will be able to cope with those context specifics by means of configuring the according applicable set of activities of type *Processing*, *Storage*, *Transferring*, *Output*, *Other* and *External*.

However, the application of the generic accounting model to a real-world Grid environment and the performed calculation exposed capabilities for further accounting model improvements as for example the consideration of load balancing aspects as well as the extension of the proposed concept of quality premiums in order to better support competition for resources. Additionally, due to the fact that detailed knowledge about the model as well as the underlying Grid infrastructure and service parameters is required, the model should be further extended in the way, that a relevant set of technical accounting and metering points respectively is defined from which relevant data can be gathered.

Finally, a sensitivity analysis considering the impact of changes with respect to modified calculation input parameters as well as scenario parameter values has been conducted. This has substantiated the identified need of a more fine-grained technical accounting based on adequate metering information.

Acknowledgements This work has been performed partially in the framework of the EU IST Network of Excellence EMANICS "Management of Internet Technologies and Complex Services" (IST-NoE-026854). Parts of this work have been funded by the German Federal Ministry of Education and Research within the D-Grid Project under contract 01 AK 800 B.

References

1. B. Aboba, J. Arkko, D. Harrington: Introduction to Accounting Management; IETF RFC 2975, October 2000, http://www.ietf.org/rfc/rfc2975.txt.
2. C. Anglano, S. Barale, L. Gaido, A. Guarise, G. Patania, R. Piro, F. Rosso, A. Werbrouk: The Distributed Grid Accounting System (DGAS); (Accessed) February 2008, http://www.to.infn.it/grid/accounting/main.html.
3. A. Barmouta: Authorization and Accounting Services for the World Wide Grid; Master thesis, University of Western Australia, June 2004, pp. 1-154.
4. A. Barmouta, R. Buyya: GridBank: A Grid Accounting Services Architecture (GASA) for Distributed Systems Sharing and Integration; 17th Annual International Parallel and Distributed Processing Symposium (IPDPS 2003), April 2003, pp. 22-26.
5. A. Beardsmore, K. Hartley, S. Hawkins, S. Laws, J. Magowan, A. Twigg: GSAX Grid Service Accounting Extensions; September 2002, pp. 1-30, http://www.doc.ic.ac.uk/~sjn5/GGF/ggf-rus-gsax-01.pdf.
6. R. Buyya, D. Abramson, J. Giddy: Nimrod/G: An Architecture for a Resource Management and Scheduling System in a Global Computational Grid; 4th International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia'2000), Beijing, China, 2000, pp. 1-7.
7. R. Byrom, J. Walk, D. Kant: User Guide for APEL – Accounting Using PBS Event Logging; March 2005, pp. 1-32, http://hepunx.rl.ac.uk/edg/wp3/documentation/apel/apel-user-guide.pdf.
8. C. Frogner, T. Mandt, S. Wethal: Cluster and grid computing: Accounting and banking systems; May 2004, http://hovedprosjekter.hig.no/v2004/data/gruppe05/files/cgc2004_report.pdf.
9. D-Grid Initiative: D-Grid Initiative; (Accessed) February 2008, http://www.d-grid.de/.
10. Daidalos: A4C Framework Design Specification; Daidalos Deliverable Update D341, September 2004.
11. G. Dreo Rodosek, M. Göhner, M. Golling, M. Kretzschmar: Towards an Accounting System for Multi-Provider Grid Environments; 10th IFIP/IEEE International Symposium on Integrated Management (IM2007), Munich, Germany, May 21-25, 2007, pp. 1-10.
12. The EGEE Project: gLite – Lightweight Middleware for Grid Computing; (Accessed) February 2008, http://glite.web.cern.ch/glite/.
13. J. Gerlach, B. Neumann, E. Moldauer, M. Argo, D. Frisby: Determining the Cost of IT Services; Communications of the ACM, Vol. 45, No. 9, September 2002, pp. 61-67.
14. The Globus Alliance: The Globus Alliance; (Accessed) February 2008, http://www.globus.org/.
15. M. Göhner, M. Waldburger, F. Gubler, G. Dreo Rodosek, B. Stiller: An Accounting Model for Dynamic Virtual Organizations; Seventh IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2007); Rio de Janeiro, Brazil, May 2007, pp. 1-8.
16. The Grid-based Application Service Provision (GRASP) Project: GRASP; (Accessed) June 2006. http://eu-grasp.net/english/default.htm.
17. F. Gubler: Accountable Units for Grid Services in Mobile Dynamic Virtual Organisations; IFI diploma thesis, University of Zurich, Zurich, Switzerland, March 2006, pp. 1-97.
18. J. H. Abawajy: Grid Accounting Service Infrastructure for Service-Oriented Grid Computing Systems, SAG 2004, pp. 168-175.
19. The Internet Engineering Task Force (IETF): Authentication, Authorization, Accounting (AAA) Charter; September 2004, http://www.ietf.org/html.charters/OLD/aaa-charter.html.
20. J. Jähnert, J. Zhou, R. L. Aguiar, V. Marques, M. Wetterwald, E. Melin, J. I. Moreno, A. Cuevas, M. Liebsch, R. Schmitz, P. Pacyna, T. Melia, P. Kurtansky, Hasan, D. Singh, S. Zander, H. J. Einsiedler, B. Stiller: The 'pure-IP' Moby Dick 4G Architecture; Computer Communications, Volume 28, No. 9, June 2005, pp. 1014-1027.
21. R. S. Kaplan, W. Bruns: Accounting and Management: A Field Study Perspective; Harvard Business School Press, December 1987.
22. R. S. Kaplan, A. A. Atkinson: Advanced Management Accounting; Third Edition, Prentice Hall, January 1998.
23. M. R. Kinney, J. Prather-Kinsey, C. A. Raiborn: Cost Accounting: Foundations & Evolutions; 6th edition, Thomson South-Western, February 2005.
24. Leibniz Supercomputing Centre (LRZ): Computing Services; (Accessed) February 2008, www.lrz-muenchen.de/services/compute/.
25. Leibniz Supercomputing Centre (LRZ): LRZ Grid Portal; (Accessed) February 2008, http://www.grid.lrz.de/en/overview.html.
26. D. Lim, Q. Thuan Ho, J. Zhang, B. Sung Lee, Y. Soon Ong: MOGAS, A Multi-Organizational Grid Accounting System; SCS International Journal of Information Technology, Vol. 11(4), 2005, pp. 84-103.
27. S. Mullen, M. Crawford, M. Lorch, D. Skow: Site Requirements for Grid Authentication, Authorization and Accounting; Global Grid Forum GFD-I.032, October 2004, http://www.ggf.org/documents/GFD.32.txt.
28. Office of Government Commerce (OGC): ITIL; June 2007, http://www.ogc.gov.uk/guidance_itil.asp.
29. T. Sandholm, P. Gardfjäll, E. Elmroth, L. Johnsson, O. Mulmo: An OGSA-Based Accounting System for Allocation Enforcement Across HPC Centers; 2nd International Conference on Service Oriented Computing, November 2004, pp. 279-288.
30. UNICORE: Distributed Computing and Data Resources; (Accessed) February 2008, http://www.unicore.eu/.