A Platform for the Integrated Management of IT Infrastructure Metrics

Christian Straube, Wolfgang Hommel, Dieter Kranzlmüller
Leibniz Supercomputing Centre, MNM-Team, Boltzmannstrasse 1, 85748 Garching
[*straube,hommel,kranzlmueller*]@*lrz.de*

***Abstract*—In this work-in-progress paper, we argue that most measurement and metrics as they are used in today's IT management do not provide a sufficient foundation for qualified upper level management decisions. By exemplary applying the state-of-the-art energy efficient metrics to SuperMUC, an energy-efficient three PetaFlop/s high performance computing system that has been put into service in June 2012 at the Leibniz Supercomputing Centre, we show that there are four major gaps between the information that can be measured on a technical level and the information that is needed for management decision making. We then present our vision of a *management cockpit* that centralizes measurement and metrics management in an organization-wide manner. It aggregates, processes and transforms metrics data into indicators for management decisions. We present the research questions, our solution approaches, and preliminary results regarding the design and implementation of this management cockpit.**

***Keywords*-IT management; measurement; metrics; governance; decision support.**

I. INTRODUCTION

For many cloud computing, high performance computing (HPC), and other large data centers, raising energy consumption costs are one of the prime motives for an in-depth examination of energy-efficient technology. Obviously, energy efficiency must be considered before investing into new hardware because, for example, CPU frequency scaling is an important functionality that helps to adjust energy consumption proportional to the current workload. However, energy efficiency is not a static property that is only relevant during purchasing decisions. On the one hand, energy-saving capabilities must constantly be monitored and controlled to ensure that they are working as expected and to keep their parameters optimized for the current workload. On the other hand, the energy efficiency of, for example, air-conditioning or hot liquid cooling systems depends on ever-changing environmental characteristics such as the current outdoor temperature.

We argue that in order to better support management decisions – such as how much money to invest into specific Information Technology (IT) infrastructure improvements – a more holistic approach to measurements and metrics is urgently required. As we show in Section II, many energy efficiency metrics have been created in the past few years. However, using SuperMUC – our three PetaFlop/s HPC system [1] that entered the top 10 of the Top 500 Supercomputer Sites list in June 2012 – as an example, without loss of generality, we demonstrate in Section III that there are four major gaps that need to be closed before metrics can directly contribute to a holistic view of IT infrastructures: At the moment, 1) the information provided by metrics is not sufficient for decisions on higher abstraction levels, 2) dependencies between metrics are not sufficiently considered, 3) organization-specific requirements cannot be incorporated adequately, and 4) there is a lack of improvement recommendations that can be deduced from the metrics values.

Our work-in-progress envisions a *management cockpit*, which centralizes measurement and metrics management organization-wide. We present its design in Section IV, using energy efficiency metrics for SuperMUC as an example of how low-level metrics can be aggregated to lay the foundation for upper-level management decisions. While works for simple metrics compositions with an immediate practical benefit already, we discuss several research questions that need to be addressed in Section V. We then outline our approach, present our preliminary results, and give an outlook to our next steps.

II. METRICS MANAGEMENT IN RELATED WORK

The related work to our management cockpit vision can be grouped into the three following categories.

Metric definition – There are a lot of metrics dealing with different aspects of energy-efficiency, for instance the measurement of the energy consumption of computing servers and clusters [2], [3], or the energy consumption in optical IP networks [4]. Further examples are TEEER [5], EPI [6], ECR, and ECRW [7], [8], [9]. All of them are defined by providing calculation and interpretation rules, partially in a very comprehensive way, but nevertheless they all focus on technical aspects of a single entity on a very low level. Hence, they do not facilitate a holistic view on the energy efficiency situation of SuperMUC, which, being a large HPC system, aggregates many different hardware components in a complex architecture.

Structuring and comparison – There is literature and ongoing research in several topics about metrics taxonomy [10], [11], classification [12], and comparison [13]. These approaches structure and compare the aforementioned metrics, but they neither aggregate different metrics to derive new statements, nor do they consider dependencies and correlations between metrics. Instead, they focus on a single specific class of metrics, like equipment-level metrics, and

Copyright (c) IARIA, 2012. ISBN: 978-1-61208-226-4

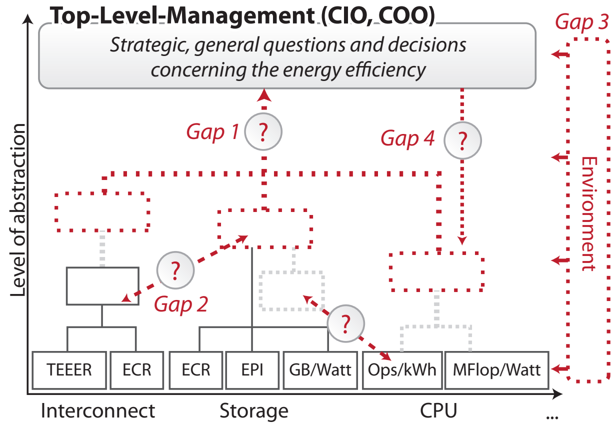

Figure 1. The four gaps in today's situation that hamper a holistic view.

confine themselves to comparison. Therefore, they do not allow a holistic view either.

Analyzing combined metrics – Besides the sole definition of metrics (group one) and the classification of those metrics (group two) there is a third group of related work that deals with the task of combining and aggregating [14] several metrics to deduce new statements about the energy efficiency of an IT infrastructure as a whole. There is no work yet that focuses on supporting upper-level management decisions. In the following, we propose our vision of a management cockpit to address this issue.

III. Problem Statement

The problem description is given with the help of the following exemplary management question: "Which components should be invested into during the next SuperMUC upgrade in order to save as much money by energy cost reduction as possible during the next n years?". By trying to answer this question, and given the previous outlined contributions of related work, we have identified four gaps, which are depicted in Figure 1. In the following explanations of the gaps, we first give a short example concerning SuperMUC and then a generalization of the problem, respectively.

Gap 1 – The Information Gap In order to answer the aforementioned management question, we first have to decide which components have the poorest energy efficiency in SuperMUC, as their potential for further improvement during the next system extension is the highest. Beside a few generally applicable metrics, most metrics can be applied only in one area, for instance an HPC/CPU metric like *MFlops/Watt* cannot be applied to storage, interconnect, or software components. Hence, there are several different metrics that have to be considered for SuperMUC as a complete system.

Generalization: For a holistic view, the (purely) technical information has only limited expressiveness and must be enriched by context and comparison information; additionally, all those information have to be aggregated to provide comprehensive information to support decision making at high level. Therefore, conversions, e. g., into currencies or hours of work, may be required.

Gap 2 – The Dependency Gap In the SuperMUC scenario, changing the CPU type to achieve a higher energy efficiency would have (strong) side effects on other components of SuperMUC. For instance, using CPUs with a smaller L2 cache size might improve the CPU energy efficiency, but at the same time, SuperMUC's system interconnect between the CPUs and the non node-local memory will have higher workloads and therefore, its energy efficiency is decreased. This may lead to a decreased overall energy efficiency.

Generalization: There are a lot of dependencies between metrics that have to be considered in order to include all relevant information. In most cases it is not adequate to improve one or only a few metrics, i. e., partial optimizations do not yield optimum results. Instead, all (involved) metrics should be improved [16], correlations have to be shown, and conclusions have to be drawn from these correlations [17].

Gap 3 – The Environment Gap Changes regarding the CPUs obviously do not only affect the storage, but also the cooling facilities. In order to assess the changed amount of energy, which the cooling facilities need for the additional CPUs, we have to compare numbers to other supercomputing centers. But we can get only useful statements if we consider the specific environment, including the location SuperMUC is deployed in: Bavaria in Germany is a relatively warm region, whereby Iceland is quite cold, so the demand for cooling would be lower there.

Generalization: The same measurements and metrics may not have the same expressiveness and purpose in different scenarios. Each organization will have to make specific adoptions to existing metrics, create complementary metrics for its specific environment, and specify, for example, how results must be interpreted properly.

Gap 4 – The Activity Gap In the end it may turn out that despite all reciprocity, using different CPUs than previously is the best way to optimize energy costs. Now we need activity recommendations that describe what to do while considering implications to other metrics and the caused costs. For instance, changing cache size has a medium energy saving effect and is quite cheap, but has a high effect on the interconnect components, whereby changing the clock speed has a high energy saving effect, is quite cheap as well, and has a low effect on the interconnect and storage components.

Generalization: In order to perform the best adoptions on the analyzed IT infrastructure, activity recommendations have to be generated out of the holistic view in a (semi-) automated way. There are a lot of challenges, such as the calculation of costs associated with each activity recommendation, and selecting the best one for a given scenario based on a consideration of the mutual reactions of metrics when changing the infrastructure, e. g., based on the application

Copyright (c) IARIA, 2012. ISBN: 978-1-61208-226-4

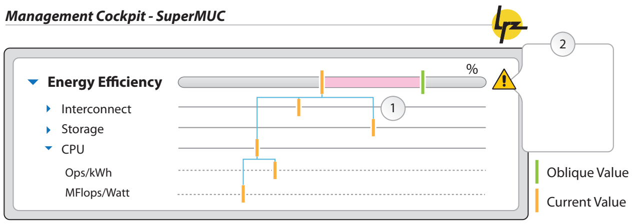

Figure 2. Our vision of a management cockpit that provides high-level, management-relevant information.

of mathematical optimization algorithms.

Closing these gaps is already a non-trivial task for the area of energy efficiency, which we use as an example here. The key challenge of our research is, however, to find a management solution that works with an arbitrary number of metrics categories and their interdependence in parallel, as we outline in the next section.

IV. Our vision of a management cockpit

To close the described four gaps, we envision a management cockpit that is built on top of an underlying comprehensive and integrated processing layer, which manages all applied metrics, and a presentation layer; a mock-up of its GUI is depicted in Figure 2. It provides an holistic view concerning the energy efficiency of SuperMUC to the top-level-management and makes all aspects feasible that were described in Section III.

Besides its main objective to "monitor, visualize and explore measurement data from different perspectives and at various levels of detail" [17] in order to characterize, evaluate, predict, and improve IT infrastructures, there are some underlying sub-objectives that shape the management cockpit, which we discuss next.

A. Considering metric dependencies

As described in Section III, metrics are currently used isolated of each other, and hence reciprocity cannot be respected. Our management cockpit shall manage all metrics that are used by an organization in a holistic way and thus make integrated statements about the SuperMUC example feasible. Besides the big advantage of respecting the reciprocity of metrics, considering metric dependencies facilitates the uncovering of strategic goal conflicts [18] and the consideration of trade-offs, for instance between energy efficiency and performance. Those trade-offs are already analyzed in some lower-level-approaches (e. g., [12]), but not yet at upper-level management.

B. Integrating measurement and metrics data

To support the holistic view and root cause analysis facilities, our solution shall integrate ("selection, embedding, and handling of the underlying data sources" [19]) and use as many data sources as required and reasonable. This leads to the problem of using and consolidating several data structures and to identify valid data contexts [17]. To be able to use the management cockpit from the first day and to avoid the "cold-start–problem" [20], existing and actual data, metrics, and measurements have to be embedded [17].

C. Using data trees

To achieve provenance for the statements provided to the management decisions, every statement must provide its data source tree to facilitate root cause analysis: starting at the top level, any aggregated metrics value can be broken down into smaller pieces and it can be explained how this high-level current value materializes. Figure 2 depicts an exemplary data tree for SuperMUC: the overall energy efficiency value is aggregated by interconnect-, storage-, and CPU-specific values, whereby the CPU value is composed of Operations/kWh and MFlops/Watt measurements.

D. Warnings and activity recommendations

As described in the aforementioned objectives, our solution shall support decision making and action planning. Therefore, there are activity recommendations that are proposed by the management cockpit, marked by (2) in Figure 2. Those recommendations depend on the delta of oblique and current values. One of the research questions we have to investigate is the modeling and creation of those recommendations; we describe this in the next section.

V. Research questions and preliminary results

To achieve our vision of the management cockpit, we have identified the following research questions, which we will analyze and answer in our future work.

A. Adequate nomenclature and classification

There are a lot of different terms that are used in the context of assessing, characterizing, and valuating the energy efficiency as well as other characteristics of an IT infrastructure, for instance *metric*, *measurement*, *quality*, *benchmark*, or *key performance indicator (KPI)*. Even if there is some literature about defining those terms (e. g., [21], [13], [22], [23], [24]) we have to define them by ourselves. Otherwise there is the risk to compare and aggregate values with different meanings and intentions, for instance a metric value and a benchmark value. To achieve this goal, we look at a metric as mathematical function and hence, it has four main components: function domain, function image, dependencies, and meta information. With this perspective, we concentrate on the characteristics of the image or range of those functions, for instance the scale, while sorting existing terms and defining our own terminology.

B. Considering metric dependencies

The area of metric dependencies is twofold: detection of dependencies, and modeling of dependencies. Both areas have individual characteristics and challenges, which have to be analyzed. The *detection of dependencies* can be done analytically or empirically. An analytical detection would look at existing information about dependencies, for instance a Configuration Management Database (CMDB) [25], and derive dependencies from those sources. An empirical detection would collect all available data at different points in time and compare them, for instance before and after a reconfiguration. After we have detected the dependencies, we have to store them in an appropriate way, so a *data model* is necessary, which must be capable of all the objectives that were introduced in the last section. Beside the aforementioned mathematical perspective, our data model divides dependencies into *reciprocity*–dependencies and *aggregation*–dependencies: the first one models correlations between metrics – for instance, improving CPU energy efficiency potentially decreases interconnect energy efficiency. The second one models the aggregation of metrics to form new statements (cf. Section IV). Interesting questions in the context of the data model are, whether there are fields that all metrics have in common, and if those fields could be placed into an abstract metrics class. This would lead to a very efficient and handy data model.

C. Derivation and generation of new statements

In the next step, we have to shape the holistic view out of the low-level approaches and thereby close gap 1. Possible solution paths are bottom-up, hypothesis generation on middle, and top-down. *Bottom-up* means that we use existing data from low abstraction levels and try to aggregate them iteratively until we reach the values that are displayed in the management cockpit. The most difficult task while doing a bottom-up generation is the "correct" selection of attributes/values at the lowest level. *Hypothesis generation* means that we formulate hypotheses on an intermediate level and try to prove or disprove those hypotheses by applying data from low abstraction levels. Those (dis)proved hypotheses are afterwards used to generate statements for the high-level management-cockpit. *Top-down* means that we start at certain points in the management cockpit and try to create the data tree beginning at the root by recursively finding suitable metrics on the next lower level. Our assumption is that we have to analyze each of those three possible ways and use a hybrid solution.

D. Target values and comparison

In order to provide "Warnings and activity recommendations" (cf. Section IV-D), target values and interpretation rules for a delta between those target values and current values are mandatory. We have to investigate how to define or rather find those target values. This step is very critical, because having wrong target values would lead to optimizing the infrastructure towards wrong values. Additionally, we have to analyze how to interpret a delta between the current value and the target value for any given metric. This interpretation has three dimensions: overall meaning, timing aspects (e. g., "delta implies the necessity to act immediately", "delta is just for the annual, paper-based report"), and impact (e. g., "the severity of the delta is very high", "solving the delta is very costly").

E. Modelling environment characteristics and their connections to metrics

The environment of an IT infrastructure can be very challenging and influences the operation and assessment of an IT infrastructure heavily. The strong impact is shown by various empirical studies, for instance [26]. To respect this fact, we want to model the environment characteristics and their connections to metrics, respectively, in order to consider this impact in the statements the management cockpits produces. The questions that arise in this context are the selection of environment characteristics that shall be modeled, the design of the data scheme to store those characteristics, and on which points those characteristics shall be connected to metrics.

VI. Conclusion and Future Work

We have shown that the energy efficiency metrics that are in use today serve their purpose of benchmarking technical components quite well, but their individual expressiveness is insufficient for IT management decisions on higher abstraction levels. Using the SuperMUC HPC system as an example, we demonstrated that the information gap, the dependency gap, the environment gap, and the activity gap need to be overcome in order to gain a holistic view on the energy efficiency of a complex IT infrastructure.

We then presented the core component of our approach, the *management cockpit*, which integrates all technical aspects of organization-wide measurement and metrics management. By aggregating, processing, and transforming existing metrics, which then are visualized for different target audiences, it is intended to be a central management tool for monitoring and decision making support. Fully overcoming the four identified gaps requires addressing several research questions first. We outlined them along with our solution approach and our preliminary results.

Our next steps include the specification of a generic metrics data model that can be applied to existing metrics and also captures their interdependence. It will serve as a basis for a prototype implementation to demonstrate the benefits of our approach in the SuperMUC real-world example. We are also working on a generalization of our approach so that besides energy efficiency, also performance, quality-of-service, and security metrics can be managed and combined in an integrated manner.

Acknowledgment

The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful discussions and valuable comments on previous versions of this paper. [27].

REFERENCES

[1] Ludger Palm, "LRZ awarded German Data Centre Prize," *Inside – Innovatives Supercomuting in Deutschland*, vol. 10, no. 1, 2012.

[2] Christos Kozyrakis and Parthasarathy Ranganathan and Suzanne Rivoire and Mehul A. Shah, "JouleSort: a Balanced Energy-Efficiency Benchmark," in *Proceedings of the International ACM Conference on Management of Data (SIGMOD '07)*, 2007.

[3] Christos Kozyrakis and Justin Meza and Parthasarathy Ranganathan and Suzanne Rivoire and Mehul A. Shah, "Models and Metrics to Enable Energy-Efficiency Optimizations," *Computer*, vol. 40, no. 12, pp. 39–48, 2007.

[4] Robert Ayre and Jayant Baliga and Kerry Hinton and Wayne V. Sorin and Rodney S. Tucker, "Energy Consumption in Optical IP Networks," *Lightwave Technology*, vol. 27, no. 13, pp. 2391–2403, 2009.

[5] L. C. Graff and T. Talbot, "Verizon NEBS TM Compliance: TEEER Metric Quantification," Tech. Rep., 2009.

[6] Sujata Banerjee and Priya Mahadevan and Parthasarathy Ranganathan and Puneet Sharma, "A Power Benchmarking Framework for Network Devices," in *Proceedings of the 8th International IFIP-TC 6 Networking Conference (NETWORKING '09)*, 2009.

[7] Luc Ceuppens and Daniel Kharitonov and Alan Sardella, "Power Saving Strategies and Technologies in Network Equipment Opportunities and Challenges, Risk and Rewards," in *Proceedings of the International IEEE Symposium on Applications and the Internet (SAINT 2008)*, 2008.

[8] A. Alimian and B. Nordman, "Network and Telecom Equipment - Energy and Performance Assessment – Test Procedure and Measurement Methodology," Tech. Rep., 2008.

[9] "Energy Efficiency for Network Equipment: Two Steps beyond Greenwashing," Juniper Networks, Inc., Tech. Rep., 2010.

[10] Ronda Henning and Ambareen Siraj and Rayford B. Vaughn, Jr., "Information Assurance Measures and Metrics - State of Practice and Proposed Taxonomy," in *Proceedings of the 36th International IEEE Hawaii Conference on System Sciences (HICSS03)*, 2003.

[11] Samee U. Khan and Lizhe Wang, "Review of Performance Metrics for Green Data Centers: a Taxonomy Study," *The Journal of Supercomputing*, 2011.

[12] Aruna Prem Bianzino and Anand Kishore Raju and Dario Rossi, "Apples-to-Apples: a Framework Analysis for Energy-Efficiency in Networks," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 3, pp. 81–85, 2010.

[13] S. C. Payne, "A Guide to Security Metrics," SANS Institute, Tech. Rep., 2006.

[14] Viktoria Firus and Ralf Reussner, "Basic and Dependent Metrics," in *Proceedings of the Dependability Metrics: Advanced Lectures [result from a Dagstuhl seminar]*, 2008.

[15] Victor R. Basili and David M. Weiss, "A Methodology for Collecting Valid Software Engineering Data," in *Proceedings of the 6th IEEE Transactions on Software Engineering*, 1984.

[16] Eduardo Fernndez-Medina and Mario Piattini and Carlos Villarrubia, "Metrics of Password Management Policy," in *Proceedings of the International Conference on Computational Science and Its Applications (ICCSA 2006), Part III*, 2006.

[17] Rudolf Ramler and Klaus Wolfmaier, "Issues and Effort in Integrating Data from Heterogeneous Software Repositories and Corporate Databases," in *Proceedings of the 2nd International ACM/IEEE Symposium on Empirical Software Engineering and Measurement (ESEM '08)*, 2008.

[18] Daniel Mares and Alessia C. Neuroni and Reinhard Riedl and Christoph Schaller and Sauter Urs, "Cockpits for Swiss Municipalities: a Web Based Instrument for Leadership," in *Proceedings of the 11th International ACM Digital Government Research Conference on Public Administration Online: Challenges and Opportunities (dg.o '10)*, 2010.

[19] Klaus R. Dittrich and Patrick Ziegler, "Three Decades of Data Integration - All Problems Solved?" in *Proceedings of the 18th World Computer Congress on Building the Information Society (IFIP)*, 2004.

[20] David M Pennock and Alexandrin Popescul and Andrew I. Schein and Lyle H. Ungar, "Methods and Metrics for Cold-Start Recommendations," in *Proceedings of the 25th International ACM Conference on Research and Development in Information Retrieval (SIGIR '02)*, 2002.

[21] Horst Lichter and Holger Schackmann, "Process Assessment by Evaluating Configuration and Change Request Management Systems," in *Proceedings of the Warm Up Workshop for ACM/IEEE ICSE 2010 (WUP '09)*, 2009.

[22] Reijo Savola, "Towards a Security Metrics Taxonomy for the Information and Communication Technology Industry," in *Proceedings of the International IEEE Conference on Software Engineering Advances (ICSEA 2007)*, 2007.

[23] John I. Alger, "On Assurance, Measures, and Metrics: Definitions and Approaches," in *Proceedings of the Workshop on Information-Security-System Rating and Ranking (WISSSR)*, 2002.

[24] "System Security Engineering - Capability Maturity Model – Model Description Document (Version 3.0)," Tech. Rep., 2003.

[25] S. Knittl, "Werkzeugunterstützung für interorganisationales IT-Service-Management - ein Referenzmodell für die Erstellung einer ioCMDB," Ph.D. dissertation, Technische Universität München (TUM).

[26] Bin Gu and Gautam Ray and Ling Xue, "Environmental Uncertainty and IT Infrastructure Governance: A Curvilinear Relationship," *Information Systems Research (INFORMS)*, vol. 22, no. 2, pp. 389–399, 2011.

[27] "Munich Network Management Team (MNM)," www.mnm-team.org, 2012.