Reinhard Brandl

Cost Accounting for Shared IT Infrastructures

GABLER EDITION WISSENSCHAFT

Reinhard Brandl

Cost Accounting for Shared IT Infrastructures

With a foreword by Prof. Dr. Martin Bichler

GABLER EDITION WISSENSCHAFT

Bibliographic information published by Die Deutsche Nationalbibliothek
Die Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data is available in the Internet at <http://dnb.d-nb.de>.

Dissertation Technische Universität München, 2007

1st Edition 2008

All rights reserved
© Betriebswirtschaftlicher Verlag Dr. Th. Gabler | GWV Fachverlage GmbH, Wiesbaden 2008

Editorial Office: Frauke Schindler / Anita Wilke

Gabler-Verlag is a company of Springer Science+Business Media.
www.gabler.de

No part of this publication may be reproduced, stored in a retrieval system or transmitted, mechanical, photocopying or otherwise without prior permission of the copyright holder.

Registered and/or industrial names, trade names, trade descriptions etc. cited in this publication are part of the law for trade-mark protection and may not be used free in any form or by any means even if this is not specifically marked.

Cover design: Regine Zimmer, Dipl.-Designerin, Frankfurt/Main
Printed on acid-free paper
Printed in Germany

ISBN 978-3-8349-0897-1

Für meine Eltern

Foreword

During the past few years, determining the "value of IT" has ranked high on the agenda of IT managers and Chief Information Officers (CIOs). The rather broad and abstract topic has been intensively discussed in the Information Systems literature for many years. It turns into a very tangible problem in the field of IT cost accounting. Nowadays, corporate information systems are distributed systems. A detailed measurement of resource demands of IT services on a distributed IT-infrastructure and respective accounting and cost allocation turns out to be very expensive and impractical in most cases. The large proportion of indirect costs and the difficulty of finding adequate allocation rates are a significant problem in practice, regularly leading to free-rider problems. This problem has largely been ignored in the academic literature so far. Dr. Brandl proposes a method to derive estimators for the resource demand of service requests in a distributed IT infrastructure. This estimator is based on a set of load tests and respective measurements as they are often performed during the deployment phase of new information systems. Cost allocation keys can now be determined based on the number of service invocations per user or per department and the respective estimators.

While these measurements provide a lean method for the determination of usage-based cost allocation keys, it is not obvious that the estimators have sufficient accuracy, in particular concerning different types of services and volatile workloads. Although it is not possible to provide tangible results for all types of information systems, Mr. Brandl performs a large number of experiments for typical multi-tiered information systems as they are in widespread use today. Queuing network models are used to validate the results for different workloads and multiple types of services. The experiments shows that resource consumption in heterogeneous environments can be predicted with high accuracy. Overall, this leads to a viable solution for the cost accounting of distributed information systems.

Dr. Brandl makes an important contribution to a largely neglected field. The book provides practical advice for IT managers for a very timely topic. I therefore hope and expect that the book will be well received not only among academics, but in particular among practitioners in cost accounting and IT controlling.

Prof. Dr. Martin Bichler

Preface

The provision of central IT infrastructure components, such as servers, storage, and networking equipment, accounts for a considerable proportion of the IT budgets of larger organizations. Typically, such components are shared among multiple applications and internal customers. Objective measurements of their respective resource consumption are technically difficult and incur high costs. In practice, infrastructure cost allocation is regularly based on simplified allocation keys which cause multiple free-rider problems and discontent among the stakeholders.

This thesis proposes a method to estimate the expected resource consumption of customer-oriented services across the components involved. The estimates are determined in a load test prior to the roll-out of an application system and then combined to so-called resource profiles. By means of these resource profiles, costs can be allocated to services or service invocations. During regular operations, consumption measurements at the different components can be omitted. The estimates therefore need to be unbiased even in cases of varying system workloads and in heterogeneous environments. Furthermore, they should support IT Capacity Planning and bridge the gap between business forecasts and IT resource planning.

The concept was implemented in a software toolkit and evaluated in a set of experiments with multi-tier database applications in a data center of the BMW Group. Queuing Network Models were used to validate the resource profiles under different system workloads. In the experiments, a surprisingly high accuracy of consumption estimates as well as of Queuing Network Model predictions could be determined. Besides the experimental validation, it was analyzed how the approach could be integrated into existing IT processes at the BMW Group.

The work presented in this thesis would not have been possible without the great support of two persons: Prof. Dr. Martin Bichler and Dr. Michael Ströbel. They gave me the

opportunity to work on this topic and provided continuous feedback, inspiration, and encouragement. I am deeply grateful to both of them.
During my time as a doctoral candidate, I was employed by the BMW Group. I greatly appreciated the pleasant working atmosphere and the kind support from all my colleagues. In particular, I would like to thank my managers Harald Raufer, Alexander Pauli, and Bernhard Huber, who gave me the freedom to pursue my research activities and provided me with interesting and responsible tasks to collect valuable business experiences far beyond the dissertation project. At the BMW Group, I was given access to all the relevant information and data center resources. However, without the technical expertise of Volker Smuda and Alexander Pochivalow, it would not have been possible to conduct more than 500 load tests during the development and evaluation phase of the software toolkit. Furthermore, I am particularly thankful to my colleagues Dr. Markus Greunz and Hartmut Liefke for their numerous comments on the viability of my approach.
I would also like to thank Prof. Dr. Helmut Krcmar who readily accepted to act as the second reviewer of the thesis and provided me with valuable advice regarding future steps. Finally, I am deeply indebted to my family for enabling me to go this way and providing continuous support, patience, and understanding throughout my whole life.

Dr. Reinhard Brandl

Contents

List of Figures

List of Tables