



Rethinking TCO: Towards a More Viable and Useful Measure of IT Costs

Joshua Greenbaum, Principal
Enterprise Applications Consulting
Spring, 2005



EAC

2303 Spaulding Avenue

Berkeley CA 94703

tel 510.540.8655

fax 510.540.7354

josh@eaconsult.com

www.eaconsult.com

Executive Summary

Despite its widespread use in the enterprise software market, *total cost of ownership* (TCO) is a troubled measure with a troubled past. Early academic studies on TCO in materials procurement, and later research by leading industry analyst firms on desktop cost of ownership, had serious flaws and problems with their accuracy and reliability. The early problems, though acknowledged by the academic research community, were largely swept under the rug as TCO emerged, particularly in the enterprise software market, as a “must-have” statistical measure for buying decisions.

The main problems with current TCO studies in enterprise software have to do with the extraordinary complexity of enterprise software implementation and on-going maintenance, and the difficulty even the best-intentioned researchers have obtaining accurate data from a large-enough sample size. A true cost of ownership analysis would have over a dozen individual factors to consider, with an extraordinary degree of variability. Issues such as product type and complexity, vertical industry, product use, implementation differences, project type, skill and experience levels, and buying power all factor into total cost of ownership.

These problems make it difficult enough to measure TCO within a single vendor’s customer base and render vendor-to-vendor comparisons virtually impossible. Enterprise Applications Consulting’s analysis of leading TCO research shows that many of these studies are plagued by a small sample size and incomplete models for total cost of ownership. These problems render the studies useless for comparative analysis in support of buying decisions, despite their widespread use for this very purpose.

While problems exist in most comparative studies, EAC believes there is a role for TCO as a means for analyzing individual vendors’ products and the cost of ownership that the vendor can offer its customers and prospects. Single-vendor TCO can potentially be used to provide comprehensive cost data to prospective customers at a level of accuracy and granularity that is not possible in typical multi-vendor studies. In addition, the vendor can use comprehensive TCO data in order to refine its product and services offerings with the goal of lowering TCO and increasing overall customer value.

SAP AG's pioneering initial efforts developing a single-vendor TCO model will yield benefits that, while not as broad as those intended by the early multi-vendor TCO studies, would be all the more valuable by virtue of their comprehensiveness and completeness. Perhaps as importantly, EAC expects that the use of TCO data by customers as well as vendors would increase significantly as the cumulative growth of the data exponentially increases a TCO database's value.

Table of Contents

Introduction: The Historical Limitations of Total Cost of Ownership Studies	page 1
The Problems with Gathering TCO Data: Comparative Challenges.....	page 3
Deconstructing Existing TCO Studies.....	page 6
Towards a Single Vendor TCO Model: Rationales and Implications	page 10
SAP's TCO Framework: An Early Evaluation of a Single-vendor TCO Effort	page 11
Conclusion: The Real TCO Metric Emerges	page 14

Introduction: The Historical Limitations of Total Cost of Ownership Studies

The notion that *total cost of ownership*, or TCO, can have a legitimate function in the analysis of enterprise spending, product value, and vendor selection_ has been around since pioneering academic research in the 1980s first identified TCO's predecessor, life cycle costing, as a valid metric in the procurement of supplies. Subsequent work in TCO, again mostly focused on materials procurement activities, firmly established the acronym as a catch-phrase for the use of long-term, comprehensive costs measurements as a management tool for qualifying purchasing decisions. For the most part these studies showed the value of TCO as a tool for looking beyond product costs in order to understand the total life-cycle costs, and thus the value, of a purchasing decision.

There were important flaws in these early studies, particularly around the ability of companies to accurately gather and understand the growing body of data required to make comprehensive TCO calculations. Furthermore, despite the availability of some relatively comprehensive TCO models for materials procurement, usage was reported to be quite low, hovering between 18 percent and 25 percent, depending on the study.

TCO's use as a tool for IT management suffered a similarly disappointing trajectory. The analyst firm Gartner Group is widely credited with introducing TCO to the IT world in its analysis of desktop computer cost of ownership in 1987. That report, and others like it, were labeled "hype" by rival analyst firm Forrester Research, which showed in a follow-up TCO study that 78 percent of IT managers don't use TCO measures in calculating their desktop management costs – a number virtually identical to usage rates for materials procurement TCO reported in the academic literature. Even though these studies, and their problems, were widely noted by observers in industry and academia, the use of TCO expanded significantly in the 1990s.

Against this background, and almost without critical evaluation, TCO emerged as a tool for comparing the relative costs of enterprise software products and implementations, and, by extension, vendors. A plethora of studies appeared from a number of industry analyst firms, and with them a mindset emerged that sought to promote TCO to the level of a "must-have" metric in the evaluation of competing enterprise software products. The result is that today major analyst

Rethinking TCO

firms like Gartner, Meta Group, and Forrester all develop and promote TCO metrics by which their end-user customers can evaluate the relative value of different vendor products.

What has been swept under the rug in the current TCO wave is that the data and usage problems of the early studies continue, even with the growing complexity of the models used to measure enterprise software TCO. This has left most TCO datasets and results hobbled by serious limitations, even as their proponents continue to promote the importance of TCO in the vendor and product evaluation process. Enterprise Applications Consulting (EAC) believes that while there is still some significant potential value to TCO studies, these limits render their use for relative evaluations between different vendors and products as incorrect and inappropriate, particularly when industry-specific and user-specific factors are taken into consideration. A new view of TCO is needed, which will provide new ways to place reasonable limits on its use as well as expand its value to prospective enterprise software customers.

This white paper will explore the issues behind the misuse of TCO and suggest that a new view of TCO be brought forth in the enterprise software market. This new view proposes that the limits of data and the complexity of any reasonable TCO models render TCO unfit for use as a relative measure of value and cost between different vendors and products. Rather, EAC proposes, TCO should be used to delineate the cost and value of implementations within the individual product families of individual vendors, based on the potential availability of much larger and more complete data sets. This single-vendor TCO model would allow TCO to be used not as a relative measure for vendor-to-vendor comparison, but as a more in-depth measure of an individual vendor's ability to provide low total cost of ownership for customers of a specific size and complexity and in specific vertical markets. This information would be of value not just to prospective buyers, but to customers looking for ways to lower existing costs by comparing their results with those of other, similar buyers. Vendors too would benefit by using a close analysis of comprehensive TCO data as a means to augment product and services development with an eye towards lowering costs and increasing value.

This white paper explores the problems with using TCO as a relative measure of IT value and cost and why a vendor-specific TCO model is a much more accurate and ultimately more useful means to evaluate enterprise software. The forthcoming TCO Framework from SAP AG, previewed by EAC, will be used as an example of this new vendor-specific TCO model.

The Problems with Gathering TCO Data: Comparative Challenges

Data quality and quantity are the fundamental problems that militate against using TCO as a comparative measure. The lack of large, comprehensive, multi-vendor datasets has rendered all major multi-vendor TCO studies reviewed by EAC as statistically invalid, and very limited in their applicability to understanding the TCO offered by a single vendor, much less as a comparative measure between vendors. Even those studies that have been based on a sound TCO model have suffered from this problem with data.

Indeed, the more comprehensive the data model for TCO, the more that incomplete data become a major problem. This underscores a fundamental limitation of most current multi-vendor TCO studies: the complexities of enterprise software and its implementation make the gathering of comprehensive data extremely difficult, if not, under most circumstances, impossible.

The breadth of functionality in enterprise software and the complexity of the implementation process significantly limit the use of TCO. While the processes that early TCO studies attempted to quantify – materials procurement, as researched by Prof. Lisa Ellram of Arizona State University, or the procurement of desktop computers, the subject of the pioneering Gartner study – are not simple to model or analyze, enterprise software is an order of magnitude even more complex. The list of variables can be extensive: differences in hardware and software costs, implementation methodology, on-going operations cost, the impact of upgrades, differential usage patterns, and potential leveragable value across the enterprise are among the many issues that make enterprise software an extremely complex product class on which to base TCO studies.

The large amount of variability in product complexity and comprehensiveness from one enterprise software vendor is an important limitation on the potential for using TCO in multi-vendor analysis. Product complexity and comprehensiveness is highly variable from vendor to vendor, and from industry to industry, depending on business and regulatory practices. The amount of functionality the software covers, and how efficiently that functionality is made available in both implementation and day-to-day use, differs widely from vendor to vendor. This makes it extremely difficult to look at the use of TCO in the enterprise software market as similar to the use of TCO for analyzing the relatively simple issues of desktop computer costs or materials procurement.

ENTERPRISE SOFTWARE TCO VARIABILITY FACTORS

- Product Type and Complexity
- Vertical Industry
- Product Use
- Implementation Differences
- Project Type
- Skill and Experience Levels: IT-IQ
- Buying Power

Further militating against the common uses for TCO in enterprise software are the problems of variability in the use of enterprise software. This variability is both broad and deep. Different industries make use of the same enterprise software product in radically different ways – witness how reporting functions vary between highly regulated industries such as pharmaceuticals and financial services and relatively unregulated industries such as electronics manufacturing. In addition, different companies within the same industry may have significantly different business processes expressed in the same software product, or may have extensively customized a standard software package in order to gain some degree of competitive advantage. Indeed, this variability in use often represents a key strategic value for the enterprise software package: by using a standard software product in a non-standard or customized fashion, many customers hope to gain competitive advantage over other companies in their industries that may be using the same or similar software. The need to differentiate through competitive business practices is often expressed by competitive differentiation in software use. The prevalence of this practice further complicates comparisons between customers in the same industries, much less customers in different industries or those using different vendor products.

The implementation process is another major factor that adds to the problems with using TCO as a vendor selection tool. Implementation costs figure as one of the largest single expenses in enterprise software, and yet they are neither standardized nor consistent from one vendor to the next or one implementer, or implementation, to the next. Much of this has to do with the different requirements of individual implementations, as well as the variability between implementations in different industries.

This renders project type a key differentiator as well: A greenfield implementation for an entirely new company, for example, will have a vastly different cost and complexity than one being done to replace a legacy system with an entirely new enterprise software product. Similarly, a highly standardized implementation with little or no customization will differ greatly from an implementation that requires a large degree of customization on top of standard software. The quality of the implementation also varies significantly – indeed, the majority of enterprise project failures and problems can be traced to quality problems with the implementation.

TCO MEASUREMENT PROBLEMS

Project Scope Differences Make Comparisons Difficult

EXAMPLES OF POTENTIALLY NON-COMPARABLE PROJECT TYPES

- Legacy Replacement
- Single Revision Upgrade
- New Architecture Upgrade
- Greenfield Implementation
- Vendor Migration Project

Variability in experience with enterprise software, at the vendor, implementer, and user level, also play a major role in the cost and scope of the implementation, and further complicates the TCO equation. This is what EAC calls IT-IQ, the cumulative knowledge and experience in enterprise software that exists across an enterprise, from the executive office on down. Companies that have been using modern Web-based enterprise software for a number of years differ greatly in their IT-IQ from those that are upgrading from a legacy, green-screen environment for the first time. Their implementation knowledge, change management skills, and user training will be vastly different from one another – all key factors in calculating an ultimate total cost of ownership. Similarly, a CIO or IT management staff that has field experience with enterprise software implementations will be more adept at every aspect of the project – scoping, vendor review and selection, consensus management, as well as the nuts and bolts of implementation – than an IT team that is confronting such a project for the first time.

Buying power and capital cost differentials are additional factors that complicate comparative TCO analysis. Labor costs, for example, represent a highly complex variable: Companies that implement enterprise software in regions such as Silicon Valley can expect higher labor costs than those that implement in the Upper Midwest, for example. But with those higher labor costs comes the potential for higher skill levels as well, which may serve to cancel out the cost differential. Buying power also differs from company to company, a result of many factors, from size and industry position to strategic value to the vendor. Both these factors lend considerable variability to the overall license cost of an implementation and the relative cost of one implementation as compared to another.

The net impact of these differences is important when looking at the problem of gathering sufficient data on which to base a legitimate multi-vendor analysis of TCO. Controlling for all these variables across even two vendors would be extremely difficult, and would require an enormous sample size in order to have enough responses in each category so as to be statistically significant. This kind of analysis – multiple linear regression – would require a significant sample size out of the universe of enterprise software users. Assuming a conservatively low total number of enterprise software user organization of 50,000, a large enough sample size to calculate simple means would require almost 400 respondents. Building a much more complex linear regression model would require hundreds more responses.

This problem of having enough data to fill out a comprehensive model of enterprise software is endemic to the use of enterprise software TCO in the market today. As we shall see below, while many of the TCO studies in the enterprise market are looking for the right kind of data in order to perform multi-vendor TCO analysis, the problems of gathering sufficient data from a sufficiently large sample across all the different variables render vendor to vendor TCO comparisons impossible.

Deconstructing Existing TCO Studies

An excellent though hardly unique example of the problem of data and sample insufficiency can be seen in an analysis of the Meta Group's widely cited TCO study, "Deriving Value from 21st Century ERP Applications," published in 2003.

Rethinking TCO

This multi-vendor study attempts to meet a number of objectives, primary among them to serve as a buyers' guide for ERP software. Its stated goals are to delineate the metrics needed to understand TCO, help customers understand costs, provide guidelines for vendor and product selection, and generally share the experiences of the user surveyed for the report.

Meta is to be lauded for the thoroughness of its intentions. The Executive Summary, available on the Meta Group website, clearly delineates an understanding of the complexities of TCO in the enterprise software space by describing the different issues it purports to look at. In the Meta study, TCO is broken down into hardware, software, professional services, and internal staff costs, and includes two years of post-implementation cost data. Meta also includes time-to-implementation and time-to-benefit analyses, offers some TCO metrics as a percent of corporate revenue, and provides a discussion of "measures of complexity," according to the Executive Summary.

Despite these laudable goals, the report fails as a comparative study of vendor TCO by virtue of the small sample size that is used for the analysis. The Executive Summary makes mention of over 200 customers, though the relative TCO chart reproduced below, once outliers have been removed, shows a total response of 172. The number of customers sampled for each vendor in Figure 1 below is well below any acceptable statistical threshold for sample size, considering all the variables needed to be covered by TCO. In an industry that measured its total customer base in the tens of thousands in 2003 – SAP alone had some 20,000 customers at the time of the Meta survey – a sample size of 172 for the vendors represented in the study, much less the small numbers for the individual vendors, is much too small on which to base the conclusions that Meta Group draws. Had Meta Group performed some standard statistical analyses, such as analysis of variance (ANOVA), and weighted the results for the differences in sample size and customer base, among other factors, a relative TCO by vendor metric could have possibly been established. (There's a high likelihood that the results would be significantly different from those in Figure 1.) Absent those basic checks and balances, the relative TCO comparisons in Figure 1 are statistically invalid, and any buying decisions based on such an analysis would be made using incorrect and misleading information.

Figure 1: Relative TCO by Vendor (Outliers Removed)

	JDE	Lawson	Oracle	PSFT	QAD	SAP
# Responses	42	25	24	26	26	29
Average	1.2%	0.48%	1.12%	0.9%	0.58%	1.36%
Median	0.83%	0.38%	0.88%	0.63%	0.58%	0.87%
Smallest	0.06%	0.05%	0.01%	0.01%	0.00%	0.05%
Largest	3.21%	1.13%	3.58%	2.65%	1.4%	5.0%

**For each vendor, the two highest relative TCO responses have been removed.*

Source: "Deriving Value from 21st Century ERP Applications,"
Executive Summary, Meta Group, 2003

Figure 2 below shows an even more egregious misuse of comparative statistics. Meta's representation of benefits from vendor to vendor is based on a total of 34 responses across six different vendors. With this extremely small sample size they cannot hope to accomplish the lofty goal of providing accurate guidelines that are representative of the individual vendors' benefits, considering the extreme variability in total customers, product complexity, average deal size, and a host of other factors.

Figure 2: Average Annual Quantified Benefits by Vendor

	JDE	Lawson	Oracle	PSFT	QAD	SAP
# Responses	12	5	3	6	3	5
Average	\$1,724,167	\$1,350,000	\$2,683,333	\$5,625,000	\$3,300,000	\$4,360,000
Median	\$800,000	\$1,000,000	\$1,000,000	\$4,500,000	\$2,800,000	\$5,500,000
Smallest	\$100,000	\$200,000	\$250,000	\$1,300,000	\$2,300,000	\$500,000
Largest	\$6,500,000	\$3,000,000	\$6,800,000	\$15,000,000	\$4,800,000	\$6,300,000

Source: "Deriving Value from 21st Century ERP Applications,"
Executive Summary, Meta Group, 2003


This is of course only one of many TCO studies that are in use in the enterprise software market. While Meta Group's study is the easiest to criticize, due to the public nature of much of its work, most TCO studies – and other end-user research reviewed by EAC – show very low sample sizes and other methodological problems that militate against their use in multi-vendor analysis.

Rethinking TCO

A final problem is the extreme variability in TCO models across the enterprise software market. The three principle analyst firms, Gartner, Forrester, and Meta, all have significantly different models that they use in the analysis of TCO in the enterprise software space. (As of this writing, Gartner Group has announced its intention to purchase Meta Group. What this will mean in terms of the newly merged company's different TCO models remains to be seen.)

These different models are capable of producing significantly different results, particularly when applied to survey sample data. SAP AG, once the main subject of enterprise software TCO research, has analyzed the differences between the principle analyst groups and found an extraordinarily wide range of variability in the results each firm offers for enterprise software TCO. (See Figure 3 below).

Figure 3: Example of costs per PC-client with 3 TCO models on the same scope of analysis



Source: SAP

This variability further confounds the issue, and offers a final compelling rationale for the abandonment of TCO as a tool for multi-vendor analysis. The stark difference between TCO values per PC-client, according to SAP's analysis, brings the main problem of TCO to the forefront: in light of these widely different results, and with the knowledge that deep data problems underlie all of these results, is it possible to use TCO for any relative measure of value in the enterprise software space?

Towards a Single Vendor TCO Model: Rationales and Implications

The twin problems of data quality and quantity, when coupled with the problems of applying low sample-size studies to multi-vendor analysis, beg the question of whether total cost of ownership can ever be used in the enterprise software space. Can the promise of TCO ever be applied in a meaningful way that can be useful for user and vendor alike?

EAC believes there is a role for TCO as a means for analyzing individual vendors' products and the cost of ownership that the vendor can offer its customers and prospects. This single vendor TCO model postulates that, while not as potentially valuable as the virtually impossible multi-vendor study, there is significant value to be gained in vendor-specific analysis of total cost of ownership. This value can be seen in two key areas: Single-vendor TCO can potentially be used to provide comprehensive cost data to prospective customers at a level of accuracy and granularity that is not possible in typical multi-vendor studies. The second value comes from the ability of the vendor to use comprehensive TCO data in order to refine its product and services offerings with the goal of lowering TCO and increasing overall customer value.

The ability of single vendor TCO to deliver this value rests in the extraordinary ability that software vendors have in gathering large survey samples from their own customer bases. EAC's research into user-based surveys by enterprise software companies shows a survey response rate that is orders of magnitude greater than can usually be achieved by vendor-neutral third parties. This is achievable by virtue of the closeness of the relationship and the potential for vendors to offer direct and indirect incentives to their customers for responding.

A vendor-driven survey effort would also tend to limit sample errors, as more rigorous sampling methodology can be applied to a more representative sample of a more refined universe of respondents, i.e. those of a single vendor. Similarly, a vendor-driven survey effort would also allow a more in-depth survey than might otherwise be possible, again based on the closeness of the relationship and the ability of vendors to offer incentives for survey participation.

A single vendor survey, based on a statistically valid sample size and a comprehensive survey questionnaire, would provide extraordinary value to the vendor and its user community. The survey results, assuming a large enough sample size, would allow for analysis along many if not

most of the complex variables that were discussed in *The Problems with Gathering TCO Data* section above. Thus, analysis by project type, complexity, implementation, usage, IT-IQ, and other key variables would be possible in a statistically meaningful way. Analysis would also be possible across different industries, and at different phases of a project. EAC expects the results of this kind of study to be dramatically different than those available from the statistically limited surveys currently in use. And when applied to the delivery of cost data to customers and for internal product and service refinement, the single vendor survey could be an extremely valuable tool for the vendor's sales and development efforts.

Some protections have to be in place in order to ensure that data quality is not compromised by a vendor-driven survey effort. Anonymity, or at least the option of anonymity, must be offered to the respondent in order to prevent survey bias. Tying the survey effort to the sales team, particularly by providing incentives to the sales team for positive results or large response rates, should also be prohibited: EAC's research into sales-driven surveys shows that inaccuracy and bias often result from the complex dynamics in the salesperson/customer dynamic. Outside audits may also be necessary, though the use of single vendor surveys for internal product refinement would militate against any attempt to tamper with or otherwise skew the data in any improper direction: such falsehoods would tend to drive development efforts in the wrong direction, with potentially damaging and costly results to the vendor itself.

SAP's TCO Framework:

An Early Evaluation of a Single-vendor TCO Effort

In mid-2004 SAP undertook a major revision of its thinking on total cost of ownership, driven by many of the problems with existing TCO studies discussed in this report. That effort led to the development of what SAP calls its TCO Framework, a comprehensive approach to a single vendor TCO model that can be applied to answer questions about total cost of ownership for customers and prospective users as well as help SAP drive its development process with an eye towards lowering costs and increasing value. EAC's review of SAP's TCO Framework shows great potential as an example of how TCO can be applied successfully in a single-vendor model to the benefit of users and vendor alike.

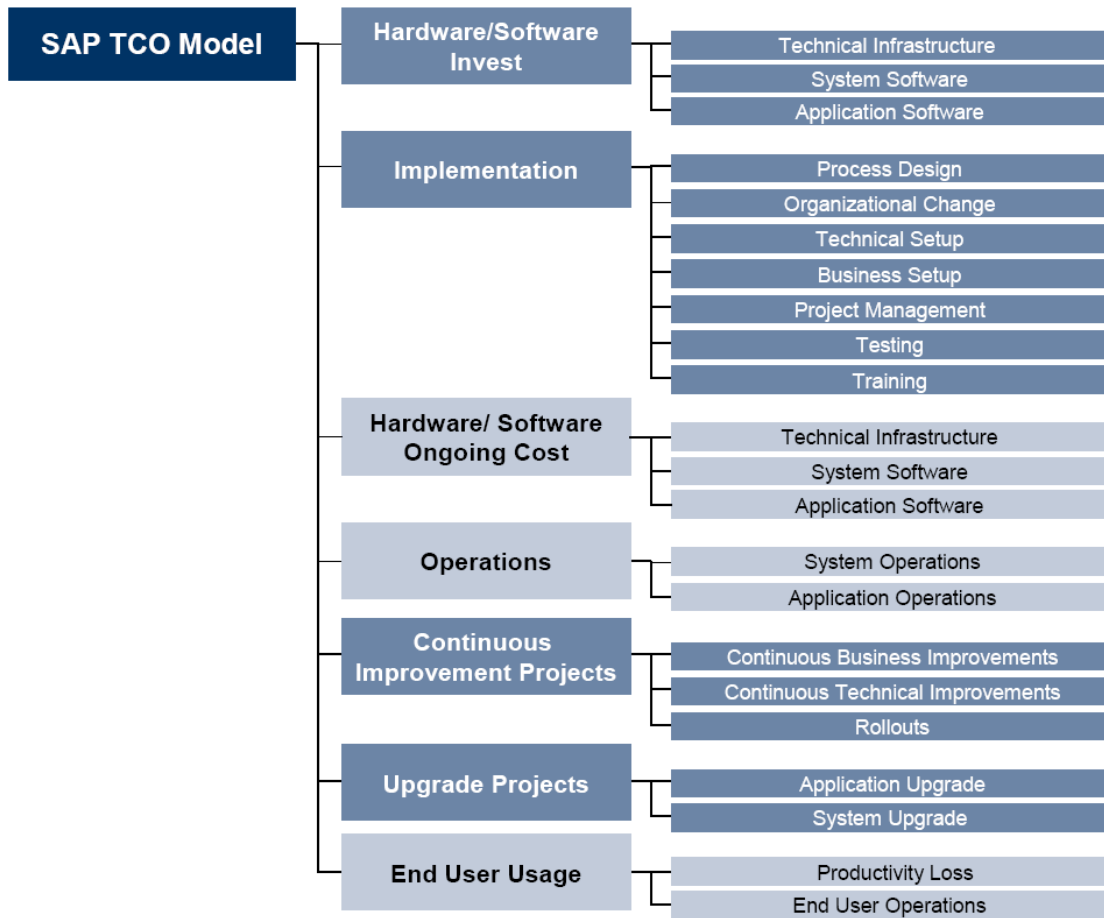
Rethinking TCO

The TCO Framework consists of five elements: a comprehensive TCO Model, Reference Parameters, a TCO questionnaire, a TCO Database, and a TCO Reduction Procedures repository. The questionnaire, derived from the TCO Model and the Reference Parameters, is used to populate the TCO Database. The TCO Reduction Procedures Repository is the accumulation of existing knowledge based on the analysis provided by the TCO Framework and other SAP resources.

The difference between the TCO Model and the Reference Parameters is important: The TCO model covers a set of issues directly related to project cost, while the Reference Parameters deal with issues related to company environment and IT practices – the context for analyzing the cost data in the TCO Model. This separation of context and cost allows SAP to better factor in issues like IT-IQ into a given analysis: by understanding how IT uses enterprise software, cost data becomes much more valuable and relevant to analysis as well as to internal developers seeking to improve products and services.

The database stores the questionnaire data and provides the platform for TCO analysis. The Database is intended to grow significantly in value over time, as the growing data set allows more complex analysis at a more granular level. SAP anticipates that initial analyses will be available to provide customers with a ranking of their implementation costs relative to other similar customers. A further analysis will provide guidelines for lowering TCO based on information in the TCO reduction procedures repository. The growing body of information in the database will create an extremely valuable asset for SAP and its customers over time. Assuming that SAP can count on a high enough participation rate among its users, this Database would provide a wealth of information of exceptional strategic value.

Figure 4: SAP's TCO Model



Source: SAP

EAC's review of the Reference Parameters and TCO Model shows that SAP is clearly on the right track in building a useful and comprehensive TCO Database. The total of 22 level-three parameters in the TCO Model cover the key cost factors in enterprise software implementations, and are comparable to or exceed the number of factors analyzed by the third party TCO studies reviewed by EAC. Similarly, the Reference Parameters cover the main issues related to the context of enterprise software implementation and use, and provide a highly useful adjunct to the cost data. Taken together, the two datasets will provide a uniquely comprehensive view of enterprise software cost and value.

It is SAP's ability to populate its TCO Database with data from hundreds of customers, however, that EAC believes will yield the greatest value to SAP and to its current and prospective

customers. This data, which by European law must be kept anonymous, will allow SAP to accomplish what the academic research and industry analyst community has only dreamt about for the last decade: a statistically valid and comprehensive analysis of enterprise software total cost of ownership. EAC believes this effort by SAP will provide a significant competitive advantage both as a sales tool as well as in SAP's development efforts.

As SAP's efforts to populate the TCO Database have only begun in early 2005, it remains to be seen how well SAP is able to fulfill the promise of its TCO Framework, but EAC expects that, as a single-vendor TCO model, SAP will be able to live up to expectations.

Conclusion: The Real TCO Metric Emerges

The low usage rates for the early TCO studies highlighted an initial problem that was perhaps reflective of their limitations. How much *total cost of ownership* in enterprise software is genuinely used as a vendor selection tool has not been well studied. With the problems outlined above, and the large differences in the results of TCO studies, however, EAC expects that an analysis of the use of TCO research utilization, particularly for vendor selection, would not necessarily reveal a large degree of success in using TCO data to make rational, effective buying decisions that would be much different than the decisions made without TCO data.

EAC believes that a vendor-specific TCO database such as SAP's, while yielding no data on comparative value between vendors, would be extremely valuable for use by the customers, prospects and developers of that vendor. The kind of comprehensive data that a vendor-driven effort can gather would yield benefits that, while not as broad as those intended by the early multi-vendor TCO studies, would be all the more valuable by virtue of their comprehensiveness and completeness. Perhaps as importantly, EAC expects that the use of TCO data by customers as well as vendors would increase significantly as the cumulative growth of the data exponentially increases a TCO database's value.

SAP's efforts in this regard bear close observation as the initiative unfolds in the coming months and years. The Framework is a well-considered effort that includes a comprehensive database and a close understanding of the relationship between cost data and the context of enterprise software use. Furthermore, EAC expects SAP to be successful in encouraging its user base to participate in

Rethinking TCO

the TCO Framework, with an increasing valuable result as user participation fills out the TCO Database.

Will TCO ever emerge as a viable multi-vendor comparison tool? One answer will come from SAP's TCO Framework: once the value of the data is known, it may spur other vendors to similar efforts. Their success in turn may provide an incentive to standardize individual vendor data sets for comparative purposes. As utopian as that may seem, at a minimum SAP's efforts will provide an important guideline for judging future TCO studies. This new effort behind a new view of TCO may herald an important revolution for IT buyers and enterprise software vendors alike.