



Comparing IBM System i Annual Operating Costs in Large Enterprises Summary of IBM Customer IT Optimization Studies 2002-2006

An IBM White Paper

EXECUTIVE SUMMARY

IBM conducted eleven enterprise IT Optimization Studies between 2002 and 2006 that included System i products. These studies examined server annual operating costs and in most cases showed marked cost and productivity differences between the IBM System i platform and alternative platforms. System i integration, reliability and virtualization technologies were frequently the reason for productivity differences. Cost differences arose from various elements including software costs, hardware and software maintenance costs, people costs and environmental costs.

IT Optimization studies examine costs to help customers understand platform efficiency. Understanding platform efficiencies provides customers a tool for selecting the optimal platform for workloads. The studies express platform efficiencies by comparing platforms' annual operating costs against workload. Server annual operating costs and staff productivity on each platform heavily influence the platform's annual operating costs. To quantify platform workload and workload efficiency uniformly across platforms, IT Optimization studies use "relative performance units" that are cross-platform, industry standard performance metrics. *Used* relative performance units indicate the amount of work the platform actually performs versus its capacity. Thus, annual operating costs (AOC) per *used* relative performance unit are a metric of platform efficiency because such analysis answers the question, "How much does it cost to perform a relative unit of work on each platform in the study?" In addition, this analysis provides a means of considering the cost implications of re-platforming applications, adding applications to a platform, consolidating a platform or implementing virtualization solutions on a platform.

Key observations include:

- System i product annual operating costs (AOC) per *used* relative performance unit averaged approximately 86% lower than Intel processor-based servers and approximately 73% lower than Unix/other midrange servers across studies with less than 1,000 servers and less than \$10B in annual revenue.
- Across all studies with de-centralized System i topologies and more than 400 servers, per server AOC averaged approximately 70% lower for System i products than for Intel processor-based servers.
- Server management staff productivity on System i products per *used* relative performance unit averaged over 4 times higher per *used* unit than on Intel processor-based servers across studies containing less than 1,000 servers.

PURPOSE AND SCOPE OF THIS PAPER

The purpose of this paper is to offer summary analysis across IT Optimization studies that included System i products. The trends and statistics cited in this paper may provide insight into general cost trends for System i products. However, because each customer environment is unique, this paper does not imply similar results will be achieved by other customers.

Because each study is customized for a client, the specific analysis differs from study to study. This paper makes note where trends occurred across clients' studies sharing a specific type of data analysis. Thus statistics cited will not represent all eleven studies as none of the studies shared 100% commonality of demographics or analysis.

BACKGROUND

Between 2002 and 2006, the IBM Systems and Technology Group's IT Optimization and Rationalization team has conducted hundreds of cross-platform customer cost studies, eleven of which included analysis of System i products.

Of the eleven studies that included System i products, four also included mainframes, Unix/other midrange servers and Intel processor-based servers. Two studies included only System i products, Unix/other midrange servers and Intel processor-based servers. Five studies included only System i products and Intel processor-based servers.

The total number of systems included in these eleven studies ranged from 62 to 5,424 with the following breakdown: 5 studies with less than 400 total systems, 4 studies with 400-999 systems and 2 studies with more than 1,000 systems. The total number of System i images in the studies ranged from 1 to 1,034 with the following breakdown: 10 studies with less than 400 System i images and 1 study with more than 1,000 images. In 5 studies, the System i topology was de-centralized or distributed while the other 6 had centralized System i machines.

The total number of Intel processor-based server images ranged from 10 to 4,337 with the breakdown as follows: 8 studies with less than 400 servers, 1 study with 400-999 images and 2 studies with more than 1,000 Intel processor-based server images. The total number of Unix/other midrange server images ranged from 4 to 1,015 with the breakdown as follows: 4 studies with less than 400 images, 1 study with 400-999 images and 1 study with more than 1,000 images. The server breakdowns are summarized in the table below:

	Studies with <400 server images	Studies with 400 to 999 server images	Studies with 1,000+ server images
Mainframe images	4	0	0
Unix/other midrange server images	4	1	1
Intel processor-based server images	8	1	2
System i images	10	0	1
Server images for all platforms	5	4	2

Figure 1. Breakdown of studies by number of server images by platform and across all platforms.

The IT Optimization studies that included System i products covered companies in the following industries: Insurance – 2, Distribution – 5, Industrial/Manufacturing – 2, Service – 1, Communications - 1. Customers ranged in size from approximately \$2B in annual revenue to over \$30B in revenue. Three studies were for customers with over \$10B per year in revenue, while the other customers had less than \$10B in annual revenue.

METHODOLOGY

IT Optimization studies examine server operating costs by analyzing numerous factors which impact cost such as lease payments, depreciation, corporate allocations, maintenance, fully-burdened staff costs, space, cooling, electrical and overhead expenses. The study follows a process by first understanding the current environment and its cost elements, analyzing data, identifying potential areas for IT optimization, developing high-level solution options, building business cases to determine future solution viability and projected costs and offering recommendations. Because of the in-depth technical and financial analysis, studies are conducted by at least one financial and one technical expert. Typical study durations are 6-10 weeks.

Financial and Technical Analysis

The studies use customers' financial, business and technical data to arrive at annual operating costs for current platforms and servers. The annual operating costs include the fully-burdened staff costs for server management, hardware and software costs, hardware and software maintenance costs and environmental costs. Additionally, studies compare relative performance capabilities across platforms and application workloads using metrics such as Ideas International's Relative Index of Performance (RIPs). These various relative performance unit metrics help size possible future systems when considering application movement to new platforms by providing a uniform metric useable across platforms. For projecting possible future solution costs, the study uses the customer's current costs and productivity ratios and applies these in conjunction with projected acquisition costs for new or upgraded environments. In projecting future solution costs, the analysis also takes into account any known migration costs.

Relative Performance Metrics

The analysis across these eleven studies uses relative performance units. Whether using TPC-C, Ideas' International's Relative Index of Performance (RIPs) or some other cross platform performance metric, studies strive to ascribe values which allow for comparison of workload and workload capabilities across platforms. Since different studies used different cross-platform metrics, this paper uses the generic term, "relative performance unit" to refer to the consolidated set of cross-platform metrics used. However, each study used the same cross-platform metric throughout that specific study. For example, if one study used the RIP metric, it was used for the entire study and no other cross-platform metric was used for that study.

Installed versus Used Relative Performance Metrics

The analysis across the eleven studies considers *installed* and *used* relative performance units. *Installed* relative performance units estimate the amount of work the server may accomplish operating at full capacity (100% utilization). *Used* relative performance units account for server utilization and thus are an estimation of the amount of work that the server actually performs. Because of very low utilization on some platforms, the analysis shows great disparity between some AOCs per *installed* and per *used* relative performance metric. In IT Optimization studies, when considering consolidated versus de-centralized topologies, virtualization implementations, or application re-platforming, the study uses the *used* relative performance units as a key indicator of workload sizing for target solutions. Examining AOC per *used* relative performance unit (r.p.u.) is one way to measure IT efficiency, specifically, “How much does it cost to perform a relative unit of work on each platform in the study?” This analysis also provides a way to consider incremental costs for moving or adding workload to various platforms.

Study Process

In understanding the current environment, the study team: 1) gathers detailed server inventory and 2) classifies servers by application, function and technical, physical and organizational boundaries. In addition to gathering extensive technical and business logistical facts, the study assembles detailed cost information about the current environment. Financial data gathered include: IT department financials, server hardware and maintenance costs for each server, software and software maintenance costs for each license, people costs and other IT costs which include environmental costs.

Data Sources

First priority in collecting financial information is to use the customer’s actual information. At times, this is not practical for the customer and instead the customer provides an estimated cost. If the customer is unable to provide an actual or estimated cost, the IT Optimization study team provides a “rule-of-thumb” value. “Rule-of-thumb” values are reviewed with and agreed to by the customer to ensure they reasonably represent the environment. These “rule-of-thumb” values come from sources such as regional or national averages or averages based upon IT Optimization Team experience. For example, if a customer does not have per kilowatt-hour costs for energy usage, the team provides a “rule-of-thumb” value based upon regional or national average rates.

Determining Averages

The analysis of this paper uses a weighted average to determine average annual operating costs and average utilization across studies. These weighted averages take into account the number of servers at a particular cost-point or utilization level.

STUDY OBJECTIVES

The objectives of an IT Optimization study typically include:

- Establishing a financial baseline, cost profile and staff productivity rate for the current state.

- Reviewing current server inventory data to understand strengths of architectures, topologies and platforms according to major server functions and applications supported.
- Identifying high value/priority solution areas for possibly reducing complexity and on-going costs
- Projecting future volumes and costs based upon current state costs and acquisition costs to establish business cases for targeted solution areas.

IT Optimization studies do not attempt to declare “the best” platform. Rather the studies try to help customers determine the optimal platform for each of their various workloads. This usually leads to recommendations for solutions on a variety of platforms.

OBSERVATIONS FROM STUDIES CONTAINING SYSTEM I PRODUCTS
Average Annual Operating Cost (AOC) per Relative Performance Unit Comparisons

In studies for customers with less than \$10B in revenue and less than 1,000 servers the annual cost per *used* relative performance unit was lower on System i products than on any other platform in the study. Across these studies, the average cost per used relative performance unit on System i products was \$32.06 per *used* relative performance unit compared to \$231.94 for Intel processor-based servers and \$122.57 for Unix/Other midrange servers. These figures contrast with the costs per *installed* relative performance units where System i products averaged \$10.41 per year per *installed* unit while Intel processor-based servers averaged \$5.96 and Unix/other midrange servers averaged \$21.14 per year per *installed* unit. Marked differences in average processor utilizations account for differences between costs per *used* performance unit and *installed* performance unit. System i products in these studies averaged 41.48% utilization, Intel processor-based servers averaged 4.07% and Unix/other midrange systems averaged 18.7% utilization.

	Average A.O.C per Installed Relative Performance Unit (r.p.u.)	Average A.O.C. per Used Relative Performance Unit (r.p.u.)	Average Utilization
Intel Processor-based Servers	\$5.96	\$231.94	4.07%
Unix/Other Midrange	\$21.14	\$122.57	18.70%
System i Products	\$10.41	\$32.06	41.48%
System i products in studies for customers with <\$10B in annual revenue and <1000 servers averaged:	174.82% higher A.O.C. than Intel processor-based servers per installed r.p.u. 50.74% lower A.O.C than Unix/Other midrange servers per installed r.p.u.	86.18% lower A.O.C. than Intel processor-based servers per used r.p.u. 73.85% lower A.O.C than Unix/other midrange servers per used r.p.u.	10.20 times higher utilization than Intel processor-based systems 2.22 times higher utilization than Unix/other midrange systems

Figure 2. Average Annual Operating Costs and Processor Utilization for studies of customers with less than \$10B in annual revenue and less than 1000 servers.

De-centralized per Server Average Annual Operating Cost (AOC) Comparisons

In studies with de-centralized System i topologies with more than 400 servers, the average Annual Operating Cost (AOC) per server for System i products was lower than the AOC for stand-alone Intel processor-based servers. Average AOC for System i products in these studies was \$3,130.19 compared to stand-alone Intel processor-based servers' average of \$10,442.50.

The stand-alone Intel processor-based server average AOC can be compared to only one de-centralized customer's study which used Integrated System x solutions. Note that an Integrated System x solution is an Intel-based server that slots into a System i product chassis to run Windows-based applications. This single study indicated an AOC of \$1,874 per server for Intel processor-based servers implemented via Integrated System x solutions.

	Average A.O.C. per server
Intel Processor-based Servers	\$10,442.50
System i Products	\$3,130.19
System i products in studies with de-centralized topologies and >400 servers averaged:	70.02% lower than Intel processor-based servers for A.O.C. per server

Figure 3. Server Average Annual Operating Costs for studies of customers with de-centralized System i topologies and more than 400 servers.

Productivity Comparisons

For studies with less than 1,000 servers the ratio of *used* relative performance units managed per full-time equivalent (FTE) staff was higher on System i products than on stand-alone Intel processor-based servers. The average for System i products was 27,262 *used* relative performance units managed per each FTE employee as compared to 5,722 on stand-alone Intel processor-based servers. For these same studies average *installed* relative performance units managed per each FTE was 77,313 for System i products and 129,761 for stand-alone Intel processor-based servers. Though staff efficiency per *installed* relative performance unit was higher on Intel processor-based servers than on System i products, Intel processor-based servers' low utilization rates indicate the staff manages a large amount of unused capacity.

Another productivity comparison came from a customer experiencing a productivity ratio of 1 FTE managing 168.3 Intel processor-based server images implemented via Integrated System x solutions. That same customer saw a productivity ratio of 1 FTE managing 20 Intel processor-based server images implemented via stand-alone servers. The average Intel processor-based productivity ratio for the other studies with stand-alone Intel processor-based servers was 18.22 server images managed per 1 FTE.

	Average Installed Relative Performance Units Managed per FTE	Average Used Relative Performance Units Managed per FTE
Stand-alone Intel Processor-based Servers	120,761	5,722
System i Products	77,313	27,262
System i products in studies with <1000 servers averaged:	36% less installed relative performance units managed per FTE than for Intel processor-based servers	4.76 times more used relative performance units managed per FTE than Intel processor- based servers

Figure 4. Productivity comparisons for studies of customers with less than 1,000 servers.

Application Trends: System i products in these studies supported core, critical business applications. In some instances they also supported key web applications, e-mail, application development or infrastructure. For the same studies, the Intel processor-based servers primarily supported infrastructure with some also supporting application development and business applications.

Solution Scenario Trends: System i customers used IT Optimization studies for varied reasons including modeling workload and server consolidation scenarios and application re-platforming. Migrating applications from System i products to other platforms always projected higher annual operating cost. Re-platforming Unix, Linux and Windows applications to System i products sometimes projected lower annual operating costs for customers. Other System i solution scenarios with projected savings included platform modernization and server and workload consolidation.

CONCLUSION

Analyzing the eleven IT Optimization studies that included System i products revealed some operating cost trends for these studies. System i products consistently had very high productivity ratios for managing servers. This provided one element towards System i products' lower annual operating costs. Ability to manage multiple workloads and environments while driving high system utilization also contributed to lower annual operating costs. The cost benefits of System i product flexibility manifested itself in many ways, often including fewer software licenses, lower software and hardware maintenance costs, lower people costs and reduced environmental requirements.

Studies projected select workloads may experience lower costs if re-platformed to System i products. Some studies included projected annual operating costs should core business applications be migrated from the System i environment. In all cases, studies projected higher annual operating costs migrating from System i products to alternative platforms.

DISCLAIMERS AND TRADEMARKS

Copyright IBM Corporation 2006

August 2006

References in this publication to IBM products or services do not imply that IBM intends to make them available in every country in which IBM operates. The information in this publication may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

IBM, the IBM logo, AIX, AIX 5L, BladeCenter, DB2, System i, System i5, i5/OS, Micro-Partitioning, POWER, System x and System p are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Novell is a registered trademark of Novell, Inc.

Red Hat is a registered trademark of Red Hat, Inc.

Other company, product, and service names may be trademarks or service marks of others. All statements regarding IBM's future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

Information concerning non-IBM products was obtained from customer studies based upon customer experiences. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

Costs for two studies were normalized to account for equivalent availability across platforms.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.