



The Reality of Rehosting: Understanding the Value of Your Mainframe

An IBM® Redbooks® Point-of-View
publication

By **Emily Farmer**
IBM Senior Analyst

Highlights

Moving applications from the mainframe to distributed environments often comes with the expectation of cost savings. However, studies reveal a conclusion that is counter to conventional wisdom:

- ▶ It could actually cost less to stay and grow on the mainframe than to move to a distributed systems environment
- ▶ Hidden costs often inflate the actual cost of moving mainframe-based applications to a distributed environment
- ▶ Workloads running on the mainframe are often a best fit for that platform, and administrators achieve significant operational efficiencies on a consolidated mainframe platform

The bottom line is that the actual cost of moving applications from mainframe to distributed environments might be greater than you think.

The first mainframe computers were introduced in the 1960s, and in the intervening years, the mainframe has become a mainstay for corporate businesses worldwide. Today, businesses trust their most mission-critical applications and data to the mainframe. Yet in recent years, some mainframe clients are attempting to move workloads off the mainframe (often referred to as *rehosting*) believing this will save them money. Typically, these clients have outdated hardware and software, smaller mainframe footprints, or perhaps a poor understanding of the true costs and value of their mainframe.

Although some service providers promote mainframe rehosting with cost savings, a careful analysis of rehosting case studies shows this claim in most cases to be untrue. Alternatively, industry trends, such as shrinking budgets, power and cooling constraints, and server sprawl, seem to favor the mainframe.

An incorrect assessment of hardware and software costs, migration costs, replacement application development costs, dual operations costs, and so on can all lead to situations where a rehosting effort can go awry. Additionally, not accounting for all the reasons why “deconsolidation” of servers occurs can result in significant cost increases. Thus, understanding the reality and true costs of rehosting workloads off the mainframe is crucial.

IBM has a team of experts (the IBM Eagle Team) to help clients understand mainframe costs and value. The Eagle Team performs customized total cost of ownership (TCO) studies to prove the cost-effectiveness of different platforms. Many of these studies show that an offload from the mainframe rarely saves money. In fact, in 96% of all mainframe rehosting cases from the IBM Eagle Team TCO studies, clients ultimately end up spending more money for an offload instead.

Rehosting costs are often underestimated for the following reasons:

- ▶ Incorrectly estimating the number of distributed servers needed to replace the mainframe
- ▶ Misunderstanding the risks and the hidden costs that are involved in a rehosting effort

Understanding your mainframe workloads and environment

Rehosting service providers tend to underestimate the number of distributed servers that are needed to replace the mainframe environment. Basic performance tests might suggest that replacing a single processor on the mainframe with a single processor core on a distributed server provides the same performance. However, simple core ratios are inherently flawed



in that the performance characteristics of both distributed and mainframe cores change with each new release of hardware.

In fact, multiple distributed cores might be needed to get the same performance as a single processor running an application in the mainframe environment. This issue is known as *core proliferation*. One resulting problem with core proliferation is that it drives up distributed software costs.

Core proliferation can happen for many reasons, but generally the reasons for this issue fall into the following areas:

- ▶ The nature of the platforms
- ▶ The nature of the workload

Consider the nature of your platform

The mainframe is a highly efficient virtualized platform that is designed to benefit from statistical multiplexing of many workloads. These workloads are colocated and run concurrently. Coupled with this design is the workload management of mixed priority workloads that is provided by the Processor Resource/System Manager (PR/SM™) firmware layer and the IBM z/OS® Workload Manager. Workload management ensures that important workloads complete at the priority that is specified.

The mainframe can run many different workloads simultaneously at high CPU utilization rates. Distributed servers are not designed for such economies of scale. Instead, they typically support fewer workloads and run at lower utilization rates to adequately provision for peak workload requirements and to meet service levels.

Thus, when workloads are moved off the mainframe, the number of distributed servers needed to support the workloads must increase. In some cases, a one-to-one mapping of applications to dedicated distributed servers is required so that each application can run efficiently and service levels can be maintained.

Similarly, non-production environments require fewer resources on the mainframe. Typically, development, test, and QA workloads run alongside production workloads but at a lower priority. Lower priority workloads might also run during off-peak hours.

IBM Eagle Team TCO studies with clients revealed that non-production environments add only about 20% of additional resources to the total mainframe costs.

Costs for a distributed platform, however, are significantly higher. Typically, development workloads are performed on a completely separate set of hardware systems from production workloads. The QA workload is also on yet another set of hardware systems, accounting for nearly 200% of additional resources to cover non-production environments on distributed platforms.

Workloads that traditionally run on mainframes are more often than not vital to the nature of the business. Most businesses have a full disaster recovery plan for mission-critical work in case of catastrophic failure. Disaster recovery on the mainframe typically makes use of features, such as spare or capacity on demand CPUs, automatic failovers, and data mirroring, which are highly reliable yet low-cost features. Conversely, disaster recovery for distributed platforms typically requires a 100% duplication of servers. The IBM Eagle Team TCO studies show that disaster recovery on the mainframe is on the order of 3% of total direct costs. By contrast, disaster recovery on distributed servers is closer to 21% on average and more than 40% for mission-critical applications.

Understand the nature of the workload

The nature of the workload is also critically important to consider when looking to a possible mainframe offload project. Consider these examples:

- ▶ System z® mainframes have dedicated processors for I/O operations. These processors are separate from the general processors used for business logic. Distributed servers use only one type of processor for both business logic and I/O operations. Workloads that have high I/O demand, such as batch processing or data-intensive transaction processing, run far more efficiently on the mainframe. Tests consistently show lower cost per workload on the mainframe for I/O intensive workloads.
- ▶ CICS/COBOL applications run far more efficiently on the mainframe than their equivalent on distributed servers. IBM has run tests using a core banking workload and has shown that the CICS/COBOL platform required 4.6 times fewer cycles per user interaction than a virtualized x86 platform. If these types of workloads are moved from the mainframe to a distributed platform, additional dedicated servers might be required to reach expected application service levels.

Estimating the distributed servers needed

When evaluating requirements for distributed servers to replace the mainframe, the Eagle Team uses “performance units” (published by independent sources) to determine the capacity for distributed servers. Performance units per mainframe MIPS become their point of comparison. Consider the following case that demonstrates core proliferation. (Although seemingly large, this case in fact yielded the smallest core proliferation ratio that the IBM Eagle Team TCO studies have encountered to date.)

A financial services company had been running a 1,660 MIPS workload on two 3-way IBM System z900 servers. This company ran production workloads as well as development, QA, and test workloads on these mainframes. Could this platform be replaced with one 8-core x86 server? No. In actuality, the migration required six 8-core x86 servers (supporting primarily front-end data ingestion) plus the addition of two 64-core IBM System p595 servers (for all the applications, data, and development). Thus, the rehosted environment would include a total of 176 distributed processors, compared to six mainframe processors, and a ratio of 482 performance units per MIPS. A 5-year TCO estimate showed that keeping the mainframes would have resulted in a cumulative spend of approximately \$17.9M, whereas moving to the distributed platform would cost \$25.4M, or a 42% increase.

You might wonder whether a crossover point for number of MIPS exists below which migration to distributed servers is always cheaper. 500 MIPS perhaps? In the studies conducted by the IBM Eagle Team, no such crossover point has been discovered. The IBM Eagle Team has been involved in tens of TCO studies in which the workload size was under 2,000 MIPS, with several cases involving workloads less than 500 MIPS.

Consider these two migration cases:

- ▶ A European bank chose to offload a small 332 MIPS (0.88 processors) workload from an IBM System z890 2-way mainframe. The workloads included test, development, and production systems, plus education, database, security, print, administration, and monitoring applications. The distributed replacement platform consisted of two 16-core servers to support most of the workload, plus four 1-core servers to support additional administration, batch scheduling, and provisioning

requirements. This configuration is a core expansion of 41 times or 670 performance units per MIPS. A 4-year TCO analysis showed that staying on the mainframe would cost \$4.9M, whereas moving to the distributed platform would cost \$17.9M, or 3.7 times more.

- ▶ A local government entity had a small 88 MIPS (0.24 processors) workload running on an IBM System z890 mainframe. The workload included only production and test applications. The distributed replacement platform consisted of four 2-core IBM System p550 servers, for an equivalence of 33 times more distributed core and 499 performance units per MIPS. A 5-year TCO study showed mainframe costs would be \$4.7M, whereas distributed platform costs would be \$8.1M, or 72% more.

Knowing the risks and hidden costs of rehosting

In addition to core proliferation, rehosting workloads from the mainframe includes a number of risks and hidden costs. Here is a small sample of some of the problems that you might encounter.

Missing functionality

A number of issues can arise from the fact that functions that exist on the mainframe do not always have a distributed counterpart. Replacement technologies are not always available for functions, such as hierarchical databases like IMS™, mainframe languages like PL/I and ASM, batch environments, 3270-style interfaces, BMS maps, VSAM file structures, and mainframe print facilities.

A couple of IBM Eagle Team TCO studies in particular highlight missing systems management functions:

- ▶ A US retail company used 200 system management products on the mainframe but only 15 of those products (7.5%) had distributed replacements. The cost of those 15 products alone was \$8.4M OTC plus \$1.8M annually.
- ▶ Another US retailer used 261 system management tools, but they could identify only 53 that had distributed equivalents. The company’s options were to either rewrite the applications to avoid the functions, generate new code to perform the functions, or perform the functions manually. Regardless of the choice, the missing systems functions added both the risk of the applications

failing to meet service levels and considerable cost to the projected TCO.

Suboptimized performance

Offloading attempts can often lead to suboptimized performance, which then can ultimately add to overall cost. A typical example of this suboptimized performance is when moving from IBM Enterprise COBOL to another COBOL.

CICS® is also a challenge to replace, as evidenced by a US Government organization that moved the CICS based vehicle registration system to a replacement distributed-based system. Prior to the migration, they were experiencing subsecond response times with the mainframe CICS application. After the migration, they experienced 30-second response times with the new application. The offload project cost \$28.3M, and it was three years late, resulting in dissatisfied customers.

“Chatty” applications that require many reads or updates with one or more databases also do not typically make good candidates to offload from the mainframe. In fact, most CICS applications access either VSAM or DB2® for z/OS data, some with very large access profiles. Although clients might deem it too risky to move data off the mainframe, they might be persuaded to move off the CICS portion. The result? What had previously been an efficient and fast “chat” between application and data now becomes constrained by TCP/IP connectivity and processing. The IBM Eagle Team TCO studies show numerous cases where this network latency has impacted desired service level requirements. Additionally, the added security risk of an exposed external network was not apparent when all components of the workload were running tightly together on the mainframe.

Risks of failed migration

Do not make a decision to go forward with a mainframe rehosting project without careful analysis, a complete understanding of the ramifications, and full knowledge of the resulting platform. Failure to complete the project can be expensive—in many ways—as illustrated in the following case studies.

Consider the case of one of the oldest property and casualty insurance operations in Canada. They made the decision several years ago to replace their mainframe and the 200 MIPS CICS/COBOL/VSAM/DB2 workloads. The vice president of information technology championed the project and was quoted as saying that the project would save the company in excess of \$1 million a year. One year later, however, the rehosting project was

abandoned because the companies hired to do the job did not have the necessary skills. The company spent millions of dollars on the conversion but had no results to show for it.

Another example of a costly failed rehosting attempt involved a large Asian bank. They set out to offload 60 MIPS of CICS/COBOL applications plus 30 MIPS of batch processing, leaving the IMS database on the mainframe. After two years, the project was abandoned because they failed to complete the extensive development necessary. The client spent \$5.7M on labor and mainframe time yet was unable to foresee an eventual completion date. No one on the 10-member team could describe the problems they encountered, and they all left the business.

What’s next: How IBM can help

If you are contemplating a mainframe rehosting effort, make sure that you examine all the factors that are involved. Go back and look again at the costs that concerned you in the first place. Examine chargeback allocations closely so that you get a full appreciation of actual costs compared to qualities of service provided. Remember that chargebacks typically reflect average costs, which will therefore underestimate distributed incremental costs (for high quality-of-service applications), while overestimating mainframe incremental costs.

Take a thorough look at the productivity of your mainframe platform compared to the equivalent distributed platform. Which generates more throughput in the least time? Which gives you the best cost per unit of work? Ensure that you are accounting for costs accurately when comparing platforms. Make sure that you are correctly accounting for all production, development, system, and environmental costs for all platforms.

Remember that a rehosting project is a large and lengthy investment. Make sure to consider all the costs (both capital and operational) and all the risks involved. What will be the return on investment and over what time period? Instead of replacing a working mainframe system, would it be wiser to make a much lower risk and lower cost investment in upgrading the mainframe? How much savings would this leave you for investment in the business instead?

When talking to companies that offer rehosting services, do not hesitate to ask for references. Make sure to get the names of their customers who have

completed similar migrations successfully. Then, go talk to these companies about their experiences.

Also, be sure to contact the IBM Eagle Team for a complete TCO evaluation (eagletco@us.ibm.com). This team can provide a personalized assessment and help you understand the true costs of your IT. They can help you understand the hidden costs and risks that might impact your decision. The Eagle TCO assessment is free of charge, but it might be the best “investment” your company makes!

Resources for more information

For more information about the concepts highlighted in the paper, see the following resources:

- ▶ IBM System z—Migrate to IBM
<http://www-03.ibm.com/systems/z/migratetoibm/compare.html>
- ▶ *The Mainframe Strikes Back: Five Reasons to Keep Your Applications on the Mainframe*, by Dr. John Shedletsky
<http://documents.bmc.com/products/documents/00/01/200001/200001.pdf>
- ▶ Follow a Multistep Process for Deciding Whether a Migration Off the Mainframe Is Warranted, 22 May 2013 ID:G00251621
<http://www.gartner.com/technology/reprints.do?id=1-1FTBDC8&ct=130530&st=sb>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: *IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.


This document, REDP-5032-00, was created or updated on June 21, 2013.



Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (or), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

CICS®
DB2®
IBM®
IMS™
PR/SM™
Redbooks®
Redbooks (logo) 
System z®
z/OS®

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.