Common planning pitfalls when offloading from a mainframe

Planning a move off your mainframe?

IT projects involving movement of workloads from one platform to another are complex, and those involving workload movement from a mainframe are no exception. Indeed, offload projects are often grossly underestimated resulting in unforeseen challenges, risk and cost.

This list of commonly observed offload planning pitfalls is based on sizing estimates from 35 IBM IT Economics studies in which clients either attempted or were considering an IBM Z® offload to a distributed environment.

The client environments varied significantly based on the volume of mainframe usage, the types of workloads they run and the size of their IT organization. Although the 35 studies spanned industries and geographies and ranged from small mainframe environments (1,000 MIPs or less) to large mainframe operations (over 50,000 MIPS), all studies identified sizing inaccuracies that would be costly for the client in the advent of a partial or full offload from the mainframe to a distributed environment.[1]

Common planning pitfalls

Not ALL impacted environments (i.e. PROD, HA, DEV, QA, DR, etc.) were considered

All 35 plans estimated the migration effort for production workloads but failed to consider some or all their peripheral environments that enable production. Non-production peripheral assets and licenses for Dev/Test, QA, DR, and HA environments can amount in aggregate to more than the production environment. When considering these additional environments the effort and cost for all the plans increased, and in many cases doubled.

Not ALL software was evaluated

Most estimates looked at the business critical applications without considering the impact of offloading automation, system management, and other software tooling required to manage the new environment.

Equivalent SW functionality was not validated in the distributed environment

Rarely did the estimates conduct a feature/function comparison of their mainframe system management software and distributed system management software equivalents. In all cases one or more products did not offer the same feature/function on all platforms, despite same or similar naming and versions. Without equivalence, additional costs needed to be factored for the purchase of additional products, or custom code development to deliver similar reporting and management functionality on the targeted distributed server environment.

IBM Z unique features and customization were not included in sizing

In all 35 cases the clients had exploited IBM Z capabilities to provide a customized environment. Over time their customization had increased, becoming a seamless part of their mainframe environment. Replication of the customization was minimized or overlooked during the requirements identification phase of their plan. To avoid costly discovery and programming efforts mid project, customization efforts had to be recalculated in order to achieve similar functionality in the new environment.

Similar to customized features are mainframe unique features that do not have a distributed equivalent, e.g. numerous Partitioned Data Sets (PDS) attributes, Security Access Facility (SAF) exits, Generation Data Group (GDG) collection and access of data, and other mainframe exclusive features. Unique automation processes, functions and output resulting from these features also needed to be evaluated for a distributed context.

Porting costs and time to completion were inaccurate

For all the IBM Z offload plans some amount of porting was required. At a minimum some COBOL/PLI needed to be re-written. Even modest porting efforts tended to be inaccurately estimated. Revised sizing costs ranged from two to ten times higher than initially estimated, and duration times were two to ten times longer than initially estimated.

Ported code stability was not considered

Most estimates assumed ported code at completion would provide the same service level as its predecessor. However, new code tends be less stable than code that has been running for 10 - 15 years. Typically several versions / releases will be required to reach comparable mainframe performance and stability. The additional effort and cost for development of successive release cycles was reevaluated to the quantify the impact of new code. Equally, if not even more important, was factoring the risk of using new code to the business (impaired function, delayed responsiveness, downtime, etc.) until the code has fully matured.

Migration effort was not accurately sized

In all cases the migration efforts were significantly underestimated. In addition to the cost of code development, these activities tended to be two to three times costlier and time consuming than estimated.

a. Install and support all new hardware
b. Install and support all new software
c. Migrate all storage
d. Migrate and test all automation scripts; include both batch and online
e. Convert IMS™ database(s)
f. Analyze and test Oracle / Db2® databases.
g. Migrate all tape data, write scripts and run books for all Job Control Language (JCL)
h. Train personnel for new environment
i. Hire additional personnel for distributed environment

Labor cost was not calculated properly

Most estimates assumed labor costs would remain unchanged. Offload planning requires an assessment of:

1. Short term impact caused by the learning curve of existing personnel (training, temporary hires, etc.)
2. Long term impact resulting from hiring and training of incremental permanent personnel.

Long term incremental labor is due to the increase in cores in distributed environments. As more cores and thus more servers are required, more resource are required to deploy and manage physical servers. Post migration, labor costs will also continue to rise as server numbers increase with new business demands.

Cost of running dual platforms / parallel environments was not considered

Almost all 35 offload plans focused on the start and end state without considering the cost during the interim. During the migration period two environments will need to be maintained, the existing mainframe environment and the future environment. This represents a significant increase in operating expense over several months or years to sustain two environments until the future environment can be declared production ready.

Impact of application code freeze during migration was not evaluated

In any IT environment code maintenance, updates, and enhancements occur regularly to meet new business requirements. In order to migrate, applications will either need to undergo a code freeze (which is usually not possible) or two versions of the application will need to be managed during the migration period. For 33 clients code freeze or dual license costs was estimated to be for three or more years. For 19 of the 33 clients using more than 5,000 MIPS, code freeze was estimated to be from five to ten years due to the extended amount of time it would take them to migrate all their applications.

Storage subsystem was not accurately sized

Most estimates assumed storage requirements would remain the same. Incremental storage hardware and software needed to be recalculated in order to compensate for Hierarchical Storage Management (HSM) efficiencies. All 35 clients concurred that their distributed storage would typically require more capacity than a z/OS® managed storage system.

Tape solution was not considered

Many mainframe environments rely on tape as their storage media. All 35 clients indicated that they were required to retain data for ten or more years. The amount of data stored on tapes varied but the effort to either convert tapes to a supported format in a distributed environment, or to use a mainframe for backup/recovery purposes needed to be included in their estimate. For many clients the tape conversion effort became one of the costliest tasks of their offload plan.

Print subsystem was not included

All 35 clients failed to do an assessment of their print requirements and their use of unique mainframe print functionality, e.g. InfoPrint®, Advanced Function Printing (AFP). Print requirements were identified and sized in a distributed context. For most, equivalent function did not exist for the targeted distributed environment, which would require additional programming efforts to be added in their plan.

High availability requirements were excluded

Distributed environments tend to have a greater need for high availability solutions. In all 35 offload plans, HA implementation costs had been overlooked. Additional hardware, software and labor for HA had to be factored into the estimate.

Security exposures were not examined

Security design in the mainframe environment may not be adequately addressed in a distributed environment. Almost all the estimates did not account for a level of security validation provided by mainframe unique components like Security Access Facility (SAF) exits. Alternative, albeit less granular, security functionality for a distributed server environment needed to be evaluated and added to the plan to control exposures. For mainframes using pervasive encryption, additional security measures and costs were assessed for a distributed environment.

Batch window requirements were not revised

Most mainframe clients leverage mainframe batch automation to its fullest. Generation Data Group (GDG) and Job Control Language (JCL) allow extensive automation including detection and resolution of job failures so batches can complete on schedule. Features like Batch pipes and Hyper-PAV do not have counterparts in the distributed environment. For all 35 clients existing batch windows were examined and redefined to conform to distributed application tools.

ETL cost savings was inaccurately calculated

Some offload plans proposed lowering costs by extracting, transforming, and loading (ETL) mainframe data to distributed servers. The ETL estimates showed a reduction in MIPS but failed to calculate the increase in distributed server count, core usage and storage as redundant data images are proliferated to multiple servers.

The ETL estimates did not factor the impact of network latency or whether their application would be a good offload candidate. Applications designed for co-located data will not handle ETL well and can consume considerable CPU resource, experience security compromises and network latency.

Additionally, the presence of multiple data copies introduces data sync issues; any changes in the source will render the copies outdated. The cost of frequent resyncs had not been considered in the planning estimates.

Coupling facility replacement not accurately sized

Mainframe coupling facilities offer more storage functionality than any other platform. Estimates often

assume this capability can be seamlessly executed across a cluster of distributed servers. Global Resource Serialization (GRS) and DFSMS™ (Storage Management System) allow the operating system to take over many of the tasks of managing storage. In a distributed environment these tasks need to be performed manually by systems programmers. For all 12 clients using coupling facilities, an equivalent solution had not been sized for the targeted distributed server environment.

Additional hardware refresh labor was not taken into consideration

Distributed HA environments recommend additional dedicated servers. Each server has to be individually decommissioned and re-provisioned. This effort is labor intensive and results in labor cost spikes each time a set of servers needs to be refreshed. All 35 planning estimates overlooked the incremental cost of distributed HA server maintenance.

Level of effort for disaster recovery was underestimated

True disaster recovery (DR) will require an implementation that mirrors all server configurations. In distributed environments disaster recovery solutions require extensive planning, testing, monitoring, and maintenance. Due to the complexity and coverage of all servers for DR efforts, labor overhead can be two to three times higher in a distributed environment. For all but three of the offload plans, estimated labor costs were underestimated and had to be resized.

Findings from client offload planning

In 34 cases reviewed, mainframe clients acknowledged that their offload project would be unsuccessful due to budget overrun, excessive time to completion and increased scope to meet unforeseen requirements. In the case of one study, the client acknowledged some sizing inaccuracies that had been calculated by its x86 vendor but was confident the project would still require only one year to complete. Three years later the migration effort concluded with significantly higher costs than initially quoted by the x86 vendor.[1]

Review of offload plans helped clients identify areas in which initial sizings had been underestimated, either by the extent of the effort or by omission of a particular task. These findings prompted the client to pause their offload planning and reevaluate the initial objective of the offload. Was it an executive strategy direction? A cost reduction initiative? In most cases the underlying issue was found to be unrelated to the platform and could actually be resolved with more effective exploitation of the mainframe.

Learn more about IT Economics

The IBM IT Economics team is a worldwide group of technical and financial consultants who work with clients to optimize their IT operations. The team focuses on identifying areas for efficiencies, cost reductions and increased business value for client business objectives.

Clients ask the team to find infrastructure and solution improvements to minimize overhead and maximize qualities of service. Areas of analysis include hardware and software purchase and maintenance costs, disaster recovery, security, datacenter costs such as networking, floorspace, energy, and labor. Analysis includes examination of projected versus actual MIPS, MLC and IPLA usage for client workloads.

Interested in knowing more about mainframe analysis? Contact the IBM IT Economics team at IT.Economics@us.ibm.com

About the author

John Gustavson, CTO of the worldwide IBM IT Economics Team, performs on-site client consulting for IT solutions. John has worked on hundreds of complex heterogenous client environments to find technical and financial efficiencies and savings, and are specialized in TCO, business value assessments, and workload placement in production, non-production, cloud and on-premises environments. John also perform analytics, hybrid cloud, IT best practice benchmarks and chargeback assessments.

[1] 35 IT Economics studies from 2012 -2018 involving review of offload projects were selected from diverse industries (40% banking, 17% central government, 14% local government, 11% insurance, and 3% each from automotive, computer sciences, consumer products, education, healthcare and retail), and different geographies (43% North America, 40% Europe, 9% Asia Pacific, 6% China, and 3% Middle East & Africa). Mainframe operations ranged in size (29% used 1,000 MIPS or less, 29% used 5,000 MIPS or less, 20% used 25,000 MIPS or less, 14% used 50,000 MIPS or less, and 9% used over 50,000 MIPS). Client workloads were comprised of IBM monthly license charges (MLC) and International Program License Agreement (IPLA) licensing and independent software vendor (ISV) licensing. Hardware was comprised of IBM Z servers running z/OS and specialty engines such as IBM z Integrated Information Processors (zIIPs). Each client engaged the IT Economics team to evaluate the workloads, the existing mainframe environment and proposed distributed environment for the offload. Of the 35 cases, five clients had already initiated IT offload activities. The other 30 were considering offload and were still in the planning phase of their project. For all 35 client studies, IT Economics consultants met on-site with the client to discuss offload planning and execution, analyzed forecasted project costs, and examined actual cost to date for those in execution mode. IT Economics analysis observed activity omissions and underestimated sizings in the offload projects and quantified offload costs for the clients. 34 clients concurred that their plans had underestimated the effort, cost and risk and halted offload project plans. One client of the 35 opted to proceed with their vendor's proposal to offload to x86 servers. The offload migration project lasted three years and cost $17M for migration, dual operating environment costs, capex and new x86 run rate over five years as projected in the IT Economics study versus the x86 vendor's migration estimate of one year.

IBM Corporation
New Orchard Road
Armonk, NY 10504
Produced in the United States of America,
12/2018