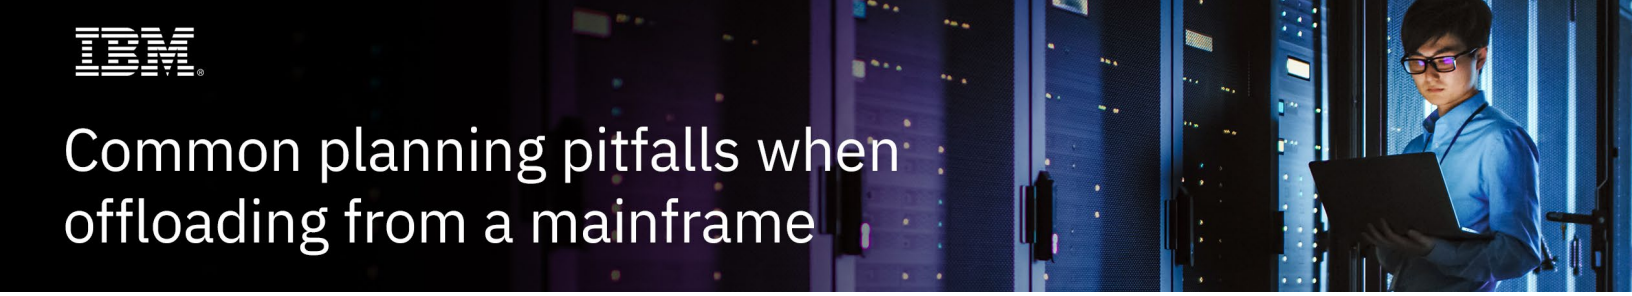

Common planning pitfalls when offloading from a mainframe

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 43 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 43 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]

Introduction

Mainframes continue to process 70% of the world's transactional workloads including critical applications for top airlines, retailers, banks, and government agencies. Organizations attempting to migrate workloads off the mainframe typically underestimate total costs by 50-75%, with large-scale projects averaging hundreds of millions of dollars in actual spending.

Mainframe modernization projects frequently face significant challenges, with 86% of organizations initiating modernization efforts but only 22% reporting success.

The Mainframe's Hidden Potential

While mainframes account for only 10–30% of IT costs, they house up to 70% of mission-critical data and workflows – like an iceberg's hidden mass beneath the surface.

Modernizing in-place with open-source solutions unlocks hidden value, accelerates innovation, and reduces operational expenses. Embedding analytics and AI within these systems can drive measurable revenue growth, delivering competitive advantage in today's rapidly evolving marketplace.

Findings from client assessments

Data from client offload analysis shows that migration is expensive and that the new run rate at project completion in many cases will be higher than a client's current IBM Z run rate.

Research from the IBM IT Economics team found a 3.2x lower annual TCO with IBM Z compared to alternative platforms.

Common planning pitfalls

Pitfall 1: 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.

Mainframe environments contain management and automation tools that need equivalent solutions in distributed or cloud environments. Where equivalents don't exist, custom replacements must be developed. Many distributed environments also require their own unique tools and skills, which migration plans often overlook.

Pitfall 2: 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.

Mainframe software provides sophisticated capabilities that often require multiple distributed products and extensive custom development to replicate. Organizations frequently discover these functionality gaps during migration, leading to additional product purchases and development work.

Pitfall 3: Integration Scale and Complexity Were Underestimated

Mainframes have far more integration points than initially documented, with hidden dependencies often leading to post-migration failures. Distributed environments struggle to handle the same volume and velocity of integrations, requiring middleware, queuing, and load balancing to compensate — adding cost, latency, and complexity.

Pitfall 4: 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 43 clients concurred that their distributed storage would typically require more capacity than a z/OS® managed storage system.

Distributed storage environments require 30% more capacity than mainframe HSM systems, with cloud storage introducing complex IOPS-based pricing that can lead to significant cost overruns. 44% of organizations frequently overprovision storage resources by 25% or more, resulting in unnecessary expenditure.

Pitfall 5: Tape Solution Was Not Considered

Many mainframe environments rely on tape as their storage media. All 43 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.

Many mainframe environments rely heavily on tape storage with 10+ year retention requirements, making tape conversion one of the costliest migration tasks. Organizations must either convert massive tape archives to supported formats or maintain mainframe systems solely for backup/recovery.

Pitfall 6: Print Subsystem Was Not Included

All 43 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.

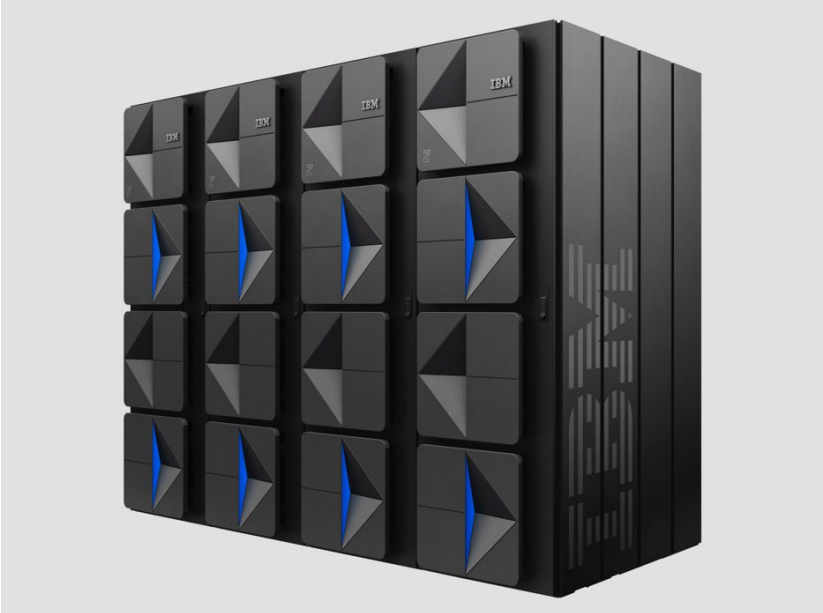

Mainframe environments utilize specialized print functionality like InfoPrint and Advanced Function Printing (AFP), which often lack direct equivalents in distributed systems. Organizations frequently discover they need extensive custom programming to replicate these capabilities.

Pitfall 7: 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 resources 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.

Distributed environments require 1.6 to 3 distributed administrators for each mainframe administrator, with short-term impacts including extensive training periods and temporary staff augmentation. Long-term labor costs continue to rise as server numbers increase, creating an ongoing multiplication of support requirements.

Pitfall 8: 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 43 planning estimates overlooked the incremental cost of distributed HA server maintenance.

Distributed environments require individual decommissioning and reprovisioning of each server, creating labor cost spikes that are 2X-3X higher than mainframe refresh cycles. While cloud environments mitigate some hardware concerns, they introduce new complexity in service provisioning and optimization.

Pitfall 9: Not ALL Impacted Environments Were Considered

All 43 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.

Pitfall 10: Coupling Facility Replacement Was 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 16 clients using coupling facilities, an equivalent solution had not been sized for the targeted distributed server environment.

Mainframe Coupling Facilities provide shared memory & synchronization, locking & coordination, and high-speed messaging ensuring low latency and high efficiency. Distributed systems rely on specialized networking (RoCE, RDMA) and software (Redis, Oracle Coherence, Apache Ignite, etc.), but this patchwork approach increases complexity, overhead, cost, and risk while failing to match mainframe reliability and performance.

Pitfall 11: Disaster Recovery Requirements Were Excluded

Distributed environments tend to have a greater need for high availability solutions. In all 43 offload plans, DR implementation costs had been overlooked. Additional hardware, software and labor for DR had to be factored into the estimate.

Distributed systems lack CBUs and GDPS® maturity, requiring overprovisioned standby infrastructure, complex failover orchestration, and costly multi-region replication. This leads to longer recovery times, higher data loss risk, and increased operational costs, while CBUs provide

backup capacity without extra software costs, making mainframes a more cost-effective DR solution.

Pitfall 12: 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 five of the offload plans, estimated labor costs were underestimated and had to be resized.

True disaster recovery in distributed environments requires mirroring all server configurations, with labor requirements typically 2X-3X higher than planned. The complexity of planning, testing, monitoring, and maintenance creates sustained operational overhead that exceeds initial estimates by 40-60%.

Pitfall 13: 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.

Even modest porting efforts consistently cost 2X to 10X more than initial estimates, with timelines extending 2X to 10X longer than planned. Assembler introduces additional complexity due to its hardware-specific dependencies, requiring extensive reengineering. These cost and timeline multipliers persist regardless of target platform, including cloud environments.

Pitfall 14: 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

Migration activities are consistently 2X to 3X more costly and time-consuming than estimated, often omitting hardware installation, storage migration, automation conversion, database migration, toolchain changes, and personnel training. Cloud migration cost overruns have become a critical challenge, with 75% of organizations substantially exceeding their budgetary projections.

Pitfall 15: 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 43 clients existing batch windows were examined and redefined to conform to distributed application tools.

Mainframe batch processing is inherently more efficient due to tightly integrated scheduling, optimized I/O, and fault-tolerant architecture. Distributed batch requires more resources, middleware, and tuning to achieve comparable performance, often at a higher cost. As a result, batch windows frequently expand, increasing processing time and delaying critical workloads.

Pitfall 16: ETL Cost Savings Were 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.

While ETL may reduce MIPS usage, it significantly impacts distributed resources, consuming 28% of total distributed cores. Every two-hour ETL cycle adds 20GB of data, resulting in operational data stores growing 2X-3X larger within one year. Network latency and security compromises further degrade performance by 15-25%.

Pitfall 17: Cost of Running Dual Platforms / Parallel Environments Was Not Considered

Almost all 43 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.

Organizations must maintain and fully fund both environments during migration, typically doubling operational costs. A $5M annual mainframe cost becomes $10M in Year 1 (dual platforms), plus $7.5M in Year 2 (with 25% savings on new platform), totaling $17.5M over the migration period. This parallel operation period often extends beyond initial estimates.

Pitfall 18: Complexity of Code Changes During Migration Was Underestimated

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 41 clients code freeze or dual license costs was estimated to be for three or more years. For 23 of the 41 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.

Organizations must manage simultaneous development across both existing and modernized codebases. Approaches like parallel development, incremental modernization, and modular refactoring demand continuous synchronization of updates, bug fixes, and new features. The effort required to maintain consistency and prevent technical debt must be factored into the offload plan.

Pitfall 19: IBM Z Unique Features and Customization Were Not Included in Sizing

In all 43 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.

Most clients have extensively customized their IBM Z environment, tailoring PDS attributes, SAF exits, and GDG collections to their specific business and technical needs. These deep customizations introduce significant blind spots in sizing assessments, turning seemingly straightforward migrations into complex and costly technological transformations.

Pitfall 20: 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 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.

New code requires multiple versions and releases to reach mainframe stability levels, with applications typically running 10-15 years on mainframe showing superior reliability. Bug-fixing costs after deployment are 10X to 100X more expensive than during initial development.

Pitfall 21: 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.

Mainframe environments possess inherent security design features, particularly through Security Access Facility (SAF) exits, that lack direct equivalents in distributed systems. Cloud environments introduce additional security challenges and network exposure risks that significantly increase implementation costs.

Pitfall 22: High Availability Requirements Were Excluded

Distributed environments tend to have a greater need for high availability solutions. In all 43 offload plans, HA implementation costs had been overlooked. Additional hardware, software and labor for HA had to be factored into the estimate.

High availability on mainframes is fault-tolerant by design, ensuring continuous uptime with self-healing hardware and OS-level resilience. Cloud and distributed environments rely on redundancy, requiring extra infrastructure, software, and operational effort. Many offload plans underestimate these costs, making HA in distributed systems more complex and expensive than expected.

Pitfall 23: Operational Excellence Practices Were Not Transferred

Mainframes achieve up to 99.999999% availability through decades of operational refinement. Migration often causes significant efficiency declines in the first year. Rebuilding stability and performance takes years, requiring major investments in infrastructure, training, and operations, increasing costs.

Findings from client offload planning

In 42 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 sizing's 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.

Conclusion

Rushing to migrate workloads off the mainframe in pursuit of perceived cost savings or modernization can lead organizations down a costly and risky path.

While public cloud and distributed platforms offer compelling benefits for certain use cases, the mainframe remains unmatched for processing high-volume, business-critical transactions that demand unwavering security, availability, and performance.

The mainframe handles business-critical workloads with superior security, scalability, and resiliency. Migration often raises costs through licensing, energy use, and added infrastructure.

Success comes not from wholesale migration, but from making informed platform decisions that align with your business objectives. For this reason, we recommend taking a fit-for-purpose approach to workload modernization and placement.

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.

For more information or to schedule a no-obligation consultation, please reach out via https://ibm.biz/ite-contact-form

About the authors

John Gustavson, formerly CTO of the worldwide IBM IT Economics Team, performed on-site client consulting for IT solutions. During his tenure, John worked on hundreds of complex heterogeneous client environments to find technical and financial efficiencies and savings, specializing in TCO, business value assessments, and workload placement in production, non-production, cloud and on-premises environments.

Linda Beckner, Senior Specialist Program Manager for IBM IT Economics, brings extensive expertise to client engagements. She works directly with organizations to optimize their IT operations, focusing on identifying cost-saving opportunities, infrastructure efficiencies, and strategic technology investments that align with business objectives.

Kathi Lafser is an IT Economics Specialist on IBM's worldwide IT Economics team. She partners with clients to uncover infrastructure efficiencies, reduce IT costs, and align technology investments with business goals through Total Cost of Ownership and Business Value Assessments. With over 20 years of industry experience, she brings particular focus to regulated industries such as healthcare, medical electronics, and automotive.

Citation Notes

[1] 43 IT Economics studies from 2012-2025 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 43 cases, eight clients had already initiated IT offload activities. The other 35 were considering offload and were still in the planning phase of their project.

For all 43 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. 42 clients concurred that their plans had underestimated the effort, cost and risk and halted offload project plans.

One client of the 43 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