

Reducing 8% Overall CPU Load for a US Insurer Without Full Infrastructure Overhaul



Overview

A US-based life insurer planned to improve its mainframe efficiency to meet growing business and customer needs. The expansion triggered higher operational, maintenance, and scaling costs. Also, transactions repeatedly triggered complex calculations, irrespective of whether the underlying data had been consuming CPU capacity without delivering additional value.

Tech Mahindra efficiently presented two targeted solutions instead of a complete infrastructure overhaul, resulting in an 8% enterprise-wide CPU reduction, capacity recovery, cost base stabilization, and enabling future growth.





Client Background and Challenges

The client is a US-based life insurer managing an extensive portfolio of operations, including investments, deferrals, pensions, and contract-to-plan processing. Its entire operations worked on the same mainframe infrastructure that processes millions of transactions daily.

The organization scaled rapidly due to increased business demand, creating a critical need to optimize its systems. However, the organization faced the following challenges:



Scaling infrastructure cost: The increase in customer and product base put the mainframe infrastructure under constant pressure. This led to rising expansion costs and a need for optimization.



Redundant computation throughout core workloads: Performing complex, data-intensive calculations on every transaction across all five workloads - whether the underlying data changed or not - creating an opportunity to eliminate significant unnecessary processing.



Excess application: Redundant application calls during retrievals within the deferral block consumed CPU resources beyond the required workload. This presented a clear opportunity for targeted optimization.


Our Solution

Tech Mahindra's approach was simple: to - eliminate processing waste at the source rather than scaling infrastructure. We deployed the following two targeted fixes simultaneously to minimize the overall CPU consumption.

(AAR) Answers at Rest: A simple solution became the core of AAR - conserving computed outputs and re-utilizing the same instead of recalculating for new calculations for every transaction. We used a methodical, phased strategy to deploy this:


- Created a dedicated Db2 performance table to store frequently generated outputs, enabling applications to reuse existing results instead of recalculating them each time.
- A publish-subscribe system was initiated to keep stored outputs up to date, automatically updating them only when required.
- Delivered the rollout in four phases over 18 months, with active monitoring at each stage to maintain production stability.
- The overall rollout was delivered in four phases over eighteen months, with active monitoring at each stage to maintain production stability.

Deferral Block Optimization: A thorough examination of deferral workloads during fetches revealed the initiation of unnecessary application calls, adding to CPU strain without serving a useful purpose. To support the broader efficiency gains from the AAR implementation, we eliminated these unnecessary calls. Thus, leading to a targeted reduction in the program-level CPU utilization within the deferral block itself.







Business and Community Impact



The AAR implementation resulted in an 8% decrease in CPU use throughout the Sys-plex. It freed up capacity and created headroom to support the ongoing expansion.



The deferral workload's program-level CPU consumption was cut by 64% by eliminating the pointless application calls.



About **Tech Mahindra**

Tech Mahindra (NSE: TECHM) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 147,000+ professionals across 90+ countries helping 1100+ clients, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal, which recognizes global companies that are actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies. For more information on how TechM can partner with you to meet your Scale at Speed™ imperatives, please visit <https://www.techmahindra.com/>.



www.techmahindra.com

www.twitter.com/tech_mahindra

www.linkedin.com/company/tech-mahindra

Copyright © Tech Mahindra Ltd 2026. All Rights Reserved.

Disclaimer: Brand names, logos, taglines, service marks, tradenames and trademarks used herein remain the property of their respective owners. Any unauthorized use or distribution of this content is strictly prohibited. The information in this document is provided on "as is" basis and Tech Mahindra Ltd. makes no representations or warranties, express or implied, as to the accuracy, completeness or reliability of the information provided in this document. This document is for general informational purposes only and is not intended to be a substitute for detailed research or professional advice and does not constitute an offer, solicitation, or recommendation to buy or sell any product, service or solution. Tech Mahindra Ltd. shall not be responsible for any loss whatsoever sustained by any person or entity by reason of access to, use of or reliance on, this material. Information in this document is subject to change without notice.