

Understanding how distributed workload drives MIPS usage

OBJECTIVE

Mainframe costs are typically driven by peak utilization, and for most customers that peak is driven by online workload coming through distributed systems. The developers of these distributed systems have limited transparency into the performance and capacity impact that their distributed programs have on the mainframe, and the mainframe operations staff have limited insight into, or ability to manage, the workload coming to the mainframe from the distributed platform. No one has a complete end-to-end understanding of the situation. Many of the tools to optimize performance and capacity on the mainframe become at best difficult to use effectively and at worst totally irrelevant.

The overall objective of this offering is the same as in the following section, 'Understanding the Mainframe 'Offload Boomerang'', to create transparency into how the distributed and mainframe components interact with each other and to provide optimization recommendations to the customer. The primary difference between the two services offerings is that this offering focuses more broadly on all workload coming to the mainframe from distributed system and not just recently offloaded workloads.

WHAT WILL WE DELIVER

Using ITBI as a starting point, SMT Data analyses the impact of the distributed workload on the mainframe. This can be done for example based on SMF 101 which gives insight into DB2's resource usage, including the DDF components. With this data it is possible to see how many MIPS are being used by remote data access, which plans or packages are involved, and which applications, servers and users are driving the MIPS usage. It is also possible to enrich this technical data with business information, for example by translating a server or user name to an organization or application.

Once the cost drivers coming from the distributed environment are identified and understood, the second step is to analyze the setup on the distributed systems based on interviews with the customer's technical staff. This includes investigation of the standard software components such as drivers. Have the correct versions been installed and have they been configured optimally? The distributed architecture and coding practices are also evaluated. Is dynamic SQL used to access the DB2 data? How does the installation ensure optimal access paths for the SQL?

The result of this offering is a presentation containing the technical findings as well as specific recommendations as to activities that can be carried out to reduce mainframe resource consumption and thereby reduce costs. Regular meetings are then held to follow up on the implementation of the recommendations, quantify cost savings, and identify new findings and recommendations.

TIME FRAME

1-3 months depending on the size and complexity of the customer's organization and cost structure.

CUSTOMER REQUIREMENTS

This offering assumes that the customer has implemented ITBI and that the basic technical data required for the analysis is already in place (e.g. SMF 101). SMT Data's consultants also require access to interview customer staff in order to understand the technical setup of both the mainframe and the distributed environment.