

MAINFRAME WORKLOAD OPTIMIZATION IS HERE!



Eliminate the Trade-off between Customer Experience and Costs

If you work in an executive position for a mainframe shop, you are likely facing simultaneous challenges to reduce costs and enhance the customer experience. That's common, but things get a lot more complicated by mature infrastructure, ever-changing impact of technology, and mainframe costs.

Does this ring a bell?

Assume your company is committed by an SLA to bill customers within three hours of a delivery of your product or service. Given the short deadline, the background workflow of jobs and applications supporting the customer billing service will be likely executed, either manually or through a workload automation tool, "as soon as possible". This means as soon as the right prerequisite conditions have been met and the right resources have been made available. But how do you ensure that workload performance does not raise MSU utilization or the LPAR 4-hour rolling average (4HRA), which ultimately determines the mainframe charge? You can't. There's no way to automate your business service workloads depending upon the dynamic load levels and therefore no way to control the MSU consumption or 4HRA peak while keeping SLAs attainment. That was true until now.

In its new [version 20](#), Control-M, the leading workload automation tool – see [Enterprise Management Associates \(EMA\) Radar Report for Workload Automation](#) – delivers the workload optimization capability for mainframe environments, to dynamically reduce the LPAR load levels, without impacting SLA performance. The solution is based on real-time measurement of load levels and policy-based workload control with a balance between costs and benefits of SLA.

This means that you will be able to establish the optimal execution of customer billing workflow, reducing MSU utilization while maintaining the three-hour SLA.

Balance Mainframe Loads while Meeting SLA

The great news is that Control-M for z/OS is now aware of real-time environments load levels, and adds this element to the orchestration logic, in combination with other criteria, like SLA requirements, prerequisite conditions and resource dependencies.

Taking into account this new element, administrators can establish sophisticated policies to run the workloads in a way that links cost control to workload SLA.

They can defer or anticipate workloads depending on their business importance, SLA impact, and system load. More jobs will be launched when the system is idle and less jobs will be launched when the system is heavily used, with a privileged pattern for SLA critical jobs. They can also reduce resource competition or waste. You will no longer have all your file transfers competing for I/O while leaving CPU idle, as you can limit the volume of file transfers to a degree which allows a balanced resource utilization among all your workloads.

Load levels are not the only element Control-M for z/OS can react to in a way that will benefit the business. Alternative examples include:

- response times of subsystems like CICS, DB2, Web server
- utilization of zIIP processors
- other elements that you can measure that may influence the mainframe processing costs.

Additionally, administrators can make decisions like:

- do not run low-urgency jobs, like internal statistics jobs, during peak 4HRA windows
- defer execution of jobs that do not jeopardize SLAs
- defer execution of jobs if CICS/DB2/Web server response time is slowing down
- defer execution of jobs that may compete to zIIP processors consumption
- limit the volume of certain jobs to not slow down online transactions
- limit the resource consumption by certain jobs to avoid competition with more important or urgent jobs
- make exceptions to the above policies when urgent jobs are putting SLA at risk.

Policies can be changed on the fly, as conditions change, so that workload optimization is always driven by current business and operational objectives, reflecting right priorities and most urgent needs.

Control-M for z/OS orchestrates workloads at appropriate times and load levels, using the appropriate resource quantities, in a way to meet those performance objectives, optimize workload execution and resource consumption, ensure consistent service level agreements and reduce mainframe licensing costs.

Conclusion

The workload optimization capability, delivered with Control-M for z/OS 20, offers a significant business advantage to mainframe executives, as it helps reducing costs while enhancing customer experience.

By integrating the dynamic load state into the workload orchestration logic, it enables you to better manage your mainframe environment, as you know how to balance your workload in a way to control MSU utilization or R4HA peaks, and therefore billing costs, and meet SLA.

Benefits include:

- optimized workload performance
- cost control linked to workload SLA
- dynamically scaled-up / scaled-down workloads
- increased efficiency
- improved resource utilization
- ensured service level agreements
- reduced costs