ThruPut Manager®

AUTOMATED WORKLOAD BALANCING

Gain further **insight** into **automated** batch workload **balancing**

This paper provides additional "under the covers" details for readers interested in how ThruPut Manager automatically balances batch workload, responding to the ever changing loading of every member in the JESplex.

Why not route every job to a specific JESplex member?

It simply takes too long. ThruPut Manager does not direct each job to execute on a specific member of the JESplex. This would be too slow, particularly if by the time that member attempts to select the job, conditions have changed and the job can't run. In that case, another member would have to be selected, notified and the whole procedure has to start again. Delays in gaining control of the JES2 checkpoint would further delay every step in the process.

Initiators would have to be available to allow the selection to occur in the first place, so central control over initiator availability would also be required. These issues all point to a complex solution with built-in delays, constantly losing precious seconds which quickly add up, particularly when typical workloads have long chains of jobs depending on the completion of one of more jobs earlier in chain.

Don't overload and spread the work around

In a nutshell, that's how ThruPut Manager automation works. Its approach can be summarized in the following points:

- On each LPAR, use the current performance of the relevant service class, LPAR and CPC when determining to select a job or not and how many initiators to have available.
- If a job can run on more than one member of the JESplex, randomize which member gets the first chance to select it by a split second.

Every 10 seconds on every LPAR on which it runs, ThruPut Manager's automation engine analyzes the performance of the CPC, LPAR and every relevant batch service class. This analysis provides the input for the ThruPut Manager algorithms that govern job selection and initiator management, automatically enabling the workload selection and balancing to respond as the load changes in this dynamic environment. The ThruPut Manager algorithms have built-in dampening to avoid wild swings, promptly responding to changes once they are more than a momentary "blip".

When load is light, chances are that all the performance indicators will permit job selection. By randomizing which member gets first chance to select new jobs, the batch workload is spread out across the available members while still honoring the system affinity of any particular batch job.

As some members of the JESplex become more heavily loaded, ThruPut Manager on those LPARs will not select some portions of the workload, naturally causing work to be run on the more lightly loaded member, making use of the available cycles.

When all members are heavily loaded, ThruPut Manager will run as much batch load on each as can receive reasonable service. When the load declines on one or more members, ThruPut Manager will select more work automatically, allowing an installation to take advantage of available capacity no matter where it is in the JESplex.

Looking back, it can be hard to see these windows of cycle availability as the SCRT only shows the peak for the month, and reports the rolling 4-hour average (R4HA), not available cycles. The RMF Post Processor reports show where there are available cycles but the resolution is limited to the size of the RMF interval, typically 15 minutes. So if there was 5 or 10 minute window when cycles were available on LPAR A, and other similar windows at different times on LPARs B, C and D, it would not be obvious from the RMF and SCRT reports.

ThruPut Manager®

THE LEADER IN BATCH AUTOMATION

8300 Woodbine Avenue, 4th Floor
Markham ON Canada L3R 9Y7
Tel: (905) 940-9404 Fax: (905) 940-5308
Email: marketing@mvssol.com www.thruputmanager.com

How does this work with soft capping?

ThruPut Manager also monitors the R4HA and can be configured to take action to reduce demand at up to 5 specified percentages of the LPAR Group Limit or LPAR Defined Capacity, called Capacity Levels, enabling lower limits to be used, thereby saving on software license fees.

With ThruPut Manager, the installation defines subsets of the batch load that are the lowest importance workloads and therefore may be held back or deferred. Among other options, ThruPut Manager allows these batch jobs to be assigned to a service class associated with each Capacity Level that has a goal type of discretionary and is a member of a WLM resource group with a system level maximum. By setting the maximum value for each resource group, the installation controls how much CPU is allocated to lower importance workload.

As with the other batch service classes, ThruPut Manager monitors these service classes and automatically controls job selection. Since these service classes are discretionary, and limited by the resource group maximum, they are likely to quickly become fully loaded causing jobs to be selected automatically on an LPAR on another CPC (and therefore not in the same LPAR Group) where the R4HA is lower and the batch is not being as constrained, if at all.

Summary

ThruPut Manager controls batch job selection and initiator deployment on each member of the JESplex. Obtaining information on service class, LPAR and CPC performance every 10 seconds, ThruPut Manager automatically runs the right amount of batch workloads, taking full advantage of available CPU cycles without overloading. ThruPut Manager ensures that workload is run wherever cycles are available resulting in an optimally and evenly loaded JESplex.

ThruPut Manager

Only certain highlights of ThruPut Manager have been discussed here. For further information, please contact us as noted on the first page.

006