

# ThruPut Manager<sup>®</sup>

AND SUB-CAPACITY PRICING

*Reduce your monthly software charges when sub-capacity pricing is used*

*This paper is of interest to datacenters using IBM's sub-capacity pricing model. It discusses the merits of managing your batch workload to lower the monthly peak in your rolling 4-hour average and consequently lower your monthly software charges. The MVS Solutions MSU Analyzer Tool is used to assess the savings in three datacenters.*

## The Key to Managing License Charges in a Sub-capacity Licensing Environment

Since charges for every license on a CPC are based on the highest Rolling 4-Hour Average (R4HA) for the month, the key to reducing your software costs is to keep the highest monthly peak as low as possible. In addition to managing costs, you must also meet the needs of your organization, so your online systems and critical batch production must not be impacted.

IBM provides soft capping as a mechanism to artificially constrain a system and therefore lower the costs. However, there is no guarantee that the constrained system won't compromise your online systems and critical batch. It is for this reason that some datacenters resist using soft capping.

We suggest that you can have the best of both worlds by deferring your low priority batch in order to lower the peaks in your R4HA. You save money without compromising the service objectives of your high importance work.

## ThruPut Manager Delivers the Solution

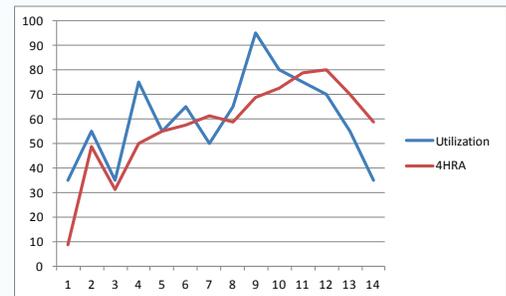
ThruPut Manager automatically monitors the rolling 4-hour average and as it approaches a peak (or the threshold you set) and defers selection of the batch workload, lowest priority first. When the peak passes it resumes selecting the deferred workload. If your peaks occur at month-end, the automated deferral of low importance work would only take place during that time.

With ThruPut Manager automation, your highest peak for the month is lowered resulting in lower monthly charges, all without affecting your online workload or higher priority batch.

## What Workloads Cause Peaks in the R4HA?

The short answer is all workloads. It is commonly thought that peaks in the R4HA are caused by spikes in online systems. While this is certainly a factor, analysis of a number of datacenters has shown that a large component of their highest peaks is batch processing. Consider that CPCs often have LPARs belonging to different JESplexes and SYSplexes, and a batch LPAR on a development JESplex can have an impact on online workload on another LPAR.

Our observation has been that a typical distribution of batch workloads can be in the region of 25% high priority, 50%



**Figure 1: When Does the Peak in the Rolling 4-Hour Average Occur?**

While standard JES2 tools can detect workload peaks (in blue on the graph), there are no tools to track the R4HA (in red on the graph). How many would predict the R4HA peak occurs at hour 12, somewhat past the workload peaks? This is when deferring low importance batch could save the datacenter many thousands of dollars per month.

# ThruPut Manager<sup>®</sup>

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](mailto:marketing@mvssol.com) [www.thruputmanager.com](http://www.thruputmanager.com)

	MSU's at peak	25% batch MSU's at peak	Savings per Month	Savings per Year
Insurance Provider (2 CPCs)	2391 MSU's	170 MSU's	\$34,000	\$408,000
Mid-sized Governmental Agency (2 CPCs)	2060 MSU's	188 MSU's	\$39,600	\$475, 200
Large Transportation Company (8 CPCs)	4826 MSU's	367 MSU's	\$73,400	\$880,800

**FIGURE 2: Potential Datacenter Savings**

*Shown here are the potential savings for three datacenters, based on the MSU Analyzer Tool and savings assumptions described earlier. As you can see, the insurance provider has higher peaks, and less of their workload is batch at those peaks, than the governmental agency. Nonetheless in both cases these savings are worth pursuing. The transportation industry company is about twice the size with about twice the savings, which suggests that the projected savings can be scaled. (In the actual reports we present the figures CPC by CPC; here the analysis is summarized.)*

medium and 25% low. So what would be the effect on your software charges of taking that low 25% of batch and, when the system is constrained and approaching its cap, deferring it for a short time until the system is less busy? Your online workload and your critical batch workloads are unaffected.

### Calculating Your Savings

MVS Solutions provides an MSU Analyzer Tool that takes data from a monthly MXG PDB and calculates the 50 highest peaks of the month and the batch component of each. We then apply the rule of thumb above by reducing each of those peaks by 25% of its batch MSUs, giving 50 new peaks. Using a typical figure of \$200/MSU/hour for the combined total of all software components that run on the CPC, the potential savings are calculated as:

$$\begin{aligned} \text{Potential Savings for 1 month} = \\ (\text{Original highest peak} - \text{New} \\ \text{highest peak}) * \$200 \end{aligned}$$

Figure 2 shows the results of this analysis for three typical datacenters.

### Assess Your Potential Savings

If you use sub-capacity pricing and would like to learn more about your potential for monthly software savings, MVS Solutions would be pleased to provide the MSU Analyzer Tool for you to run at no charge. You will better understand your monthly peaks and the benefits of using **ThruPut Manager** in your datacenter.