Sizing iSeries Disk with IBM Workload Estimator
Version 1.1

Keith Zblewski – IBM Workload Estimator Project Leader
January 20, 2006

Introduction

The IBM Workload Estimator (WLE) has been available since 1999 to size AS/400 and iSeries systems. WLE has the capability to size processor, memory, and disk resources, but the support for sizing disk had some limitations in the early versions of the tool. Over the past couple of years, several improvements have been made to WLE for sizing disk resources. This document provides an update on the current capability in WLE for sizing disk and provides some background about the methodology that WLE uses in the sizing process.

Getting Started – Setting the Target Disk Configuration

Before WLE can determine the recommended disk sizing, it needs to know which Disk I/O Adapters (IOAs) are being proposed and the speed of the proposed drives. This can be specified in the WLE Options file for iSeries which is shown below.

The key information in the Options file that is used for sizing the new disk solution includes:

- The level of disk protection - RAID-5, Mirrored, or No Protection can be specified. RAID-6 will be supported in WLE release 2006.1.
- Disk Storage Percent Full - between 0-100% with 85% being the default
- Disk Attachment Type – the following IOAs can be specified
    - 2780, 2757, 5703, or 5709
    - Original IOA/IOP – used for IOAs not specified above

- Disk Storage Type – 7,200, 10,000, or 15,000 rpm can be specified
- Target Disk Busy Rate – between 5-40% with 25% being the default [1]

Once the target configuration is specified, you can begin the sizing process for upgrades, server consolidations or new workloads.

Sizing Disk for Upgrades and Server Consolidations

There are two primary ways to size upgrades and server consolidations with WLE. One way is by using the Existing workload. The other way is by exporting performance data from PM iSeries into WLE.

Users of the Existing workload have the ability to enter the key information for the existing system that is being upgraded or consolidated. This information includes the Drive Attachment (IOA), the Drive Type (Speed), the type of data protection being used, the number of drives, and the total amount of storage used in gigabytes. Up to 8 disk groups can be specified in an Existing workload which allows you to size systems with up to 8 different types of IOA/drive speed/data protection combinations. When exporting performance data from PM iSeries, this disk information is obtained from the customer's performance and configuration data and passed to WLE.

Users of the Existing workload also have the ability to specify the number of read and write operations per second and the average size of the read and write operations for each disk group. The I/O operation frequency and size are used when calculating the change in the disk capacity required if switching from RAID-5 to Mirroring, or performing some other type of data protection change. More information on this topic is found in the section of this paper titled, "Sizing Mirrored Environments".

WLE calculates the required number of disk drives needed for the new system based on performance measurements taken for each type of disk by the IBM performance team. This team measures the capability of each drive based on the number of I/O operations the drive can perform at a target disk busy rate. WLE assumes a maximum disk busy rate

[1] Target Disk Busy Rate is being implemented in release 2006.1 of WLE, but is not currently available at the time this document was written.

of 25% when sizing all workloads, but this rate can be reduced by changing the target rate in the WLE Options file as described earlier.

WLE uses an internal disk capacity table which contains the capacity of each IOA, drive speed, and disk protection type based on the measurements described above. If the existing system being upgraded or consolidated contains a different disk configuration than will be used on the new system (as specified in the WLE options file), adjustments to the recommendation will be made to account for this difference. For example, if the existing system contains 2757 IOAs with 10k rpm drives and RAID-5 disk protection, the measurements from the IBM performance team suggest that each drive can perform approximately 53 I/O operations/second with good performance. If the new system you are upgrading to contains 2780 IOAs with 15k rpm drives and RAID-5 disk protection, each drive is estimated to perform approximately 83 I/O operations/second with the same level of performance. Thus, the new system requires about 64% of the disk arms required on the old system (53/83) and WLE will size the new system accordingly.

In a future release, the disk arm busy rate will be an additional factor used to determine the number of disks that are required. For example, if the target arm busy rate in the options file is set to 20%, but the old system has an average arm busy rate of 25%, then more disk arms will be estimated in order to reduce the busy rate to the specified target.

Sizing Disk for New Workloads

When new workloads are added to a new or existing iSeries system, WLE can be used to estimate the number of disk arms required to support the new workload. Each workload supported in WLE such as Traditional OLTP, Domino, WebSphere, etc. now estimates the number of disk operations for a given workload definition. Measurements taken by the POWER Systems Performance team determine an estimated number of read and write operations for a given number of users, transactions, etc. that are entered for each workload. Then the same capacity table described in the previous section is used to determine the number of arms that are needed for each workload.

The read and write operations estimated for each workload assumes that no disk protection is being used. When the user of WLE indicates that RAID or mirrored protection is being used, the number of read and write operations are adjusted accordingly. For example, when RAID-5 protection is used, one additional write and two additional reads are performed by the disk subsystem for each “system” write operation that occurs. Thus, the number of estimated operations are increased accordingly to reflect this behavior when RAID-5 is used.

When mirrored protection is used, one additional write operation is performed for each “system” write operation so the number of estimated write operations is doubled from its measured rate. Because the internal disk capacity tables are based on I/O rates, the adjustments made for RAID and mirrored protection cause the correct number of disk arms to be estimated in these environments.

Sizing Mirrored Environments

When sizing upgrades or server consolidations with the Existing workload, no additional adjustments are made by WLE if both the old and new system are using mirroring. However, when moving from no disk protection to a mirrored environment, the number of write operations input in this workload are multiplied by 2. Then the internal disk capacity tables are used to determine the number of disk arms required for this workload.

Likewise, if sizing a new workload with mirrored protection, the number of write operations estimated for the workload are multiplied by 2. Then the same disk capacity tables are used to determine the required number of disk arms.

Summary

The recent and upcoming improvements in WLE to size disk resources will help you do a better job when dealing with a mix of IOA types, drive speeds, and data protection options. For more help with WLE, use the Help/Tutorials tab in the main dialog window of the tool which is located at http://www.ibm.com/eserver/iseries/support/estimator.