IBM @server i5 520 Disk Performance Considerations #5709/5703 & #6574/84/94

June 16, 2004

There has been some discussion recently of how to best configure a model 520 base system or CEC (Central Electronics Complex) for disk performance. This paper provides tips and guidelines for configuring disk in a model 520 CEC.

Many customers with a small number of disk drives tend not to stress their server too heavily. This discussion will generally not be of interest to them. But for other customers who work their server harder, especially those with high disk drive utilization, this information may be of interest. Some customers will be quite satisfied with #5709 performance driving 6-8 disk drives. Other customers who use their server differently will prefer only 3-4 disk drives driven by their #5709.

It is fairly common to see iSeries disk usage averaging 25-30 disk ops per second. A reasonably busy 15k rpm disk drive can average around 60 disk ops per second. A very busy 15k rpm disk drive can do many more operations – say 100 or so. A 10k rpm disk drive is very busy at 60 ops per second.

Rule of thumb (your mileage may vary): The embedded #5709 RAID Enabler performance is expected to be about the same as, or perhaps slightly less than, the #5703 disk controller performance for most customers. The #5709 only handles a maximum of eight disk drives while the 5703 handles up to twelve drives, so its capabilities should be a good match for many customers. Initial test runs have indicated the #5709 can handle about 320 disk operations per second with minimal bottle-necking caused by the disk controller. Based on this, a good sizing rule of thumb for the #5709 would appear to be "240-280 disk operations per second".

Consider changing your single #5709 controller configuration if you have:

- three or more 15k rpm disk drives which have extremely high utilization or
- four or five or more 15k rpm disk drives which are reasonably busy or
- four or five or more 10k rpm disk drives which are very busy,

Then you will probably want to modify your configuration to add a disk controller by:

- Moving some disk from the CEC to an I/O tower/drawer like the #5095/0595/5094. There, the disk can be driven by either a #5703 or a #2757 disk controller.
- Adding a #5703 disk controller to the 520 CEC and drive any disk in the second set of four disk slots in the CEC

Additional Background -- #6574, #6584, #6594

In the 520 CEC there is an embedded disk controller with no RAID capability and no write cache. There is also an optional card (#5709) which goes into a special slot (sort of a daughter card concept) with RAID capability and 16MB write cache. Either of these disk controllers can drive up to 8 disk drives in the CEC. Neither of these controllers uses a PCI-X slot and will be referred to as embedded controllers.

There are 8 disk slots in the 520 CEC. Four are always enabled and with #6574 you enable the remaining 4 disk slots. All eight slots are then driven by the embedded disk controller. If you use the #6594 or #6584 feature code you must drive the remaining 4 disk drives with a disk controller in a PCI-X slot in the 520 CEC. The first four disk slots are always driven by an embedded controller.

If using the #6584, either the #5715 (non-RAID) or #5703 (RAID) disk controllers can drive the #6584 enabled disk. No other disk controller is supported. And because of what PCI-X slots are available for the disk controller with #6584, the disk controller **can NOT** be supported by a second IOP. This means that you cannot have a second i5/OS partition in just the CEC. You could have a partition where the disk controller was owned by Linux or AIX. You could have the disk controller owned by the same i5/OS as owns the embedded disk controller. You can have other i5/OS partitions using disk in I/O towers.

With the #6594, you have more flexibility. Either the #5715 (non-raid) or #5703 (raid) disk controllers can drive the #6594 enabled disk. No other disk controller is supported. And because of what PCI-X slots are available for the disk controller with #6594, the disk controller **CAN** optionally be supported by a second IOP. This means that you can have a second i5/OS partition in just the CEC. You could have a partition where the disk controller was owned by Linux or AIX. You could have the disk controller owned by the same i5/OS as owns the embedded disk controller. #6594 currently has 4Q04 availability. We are working to make it available 3Q04. Given the #6594 is a superset of the #6584's functionality; expect the #6584 to be withdrawn when #6594 is available.

The prices of #6594 and #6584 are the same, $265 (USA list). This means that it is a relatively low cost change to purchase a #6594/6584 to replace an installed #6574 if additional disk drives or changing disk usage indicates a change. For example, a customer could start with five disk drives and a #6574, and then replace the #6574 with a #6594 with another controller. A sixth disk drive would be required if RAID was being used by both disk controllers as 3 drives are required at a minimum for a RAID set.

#2757 Observation

There are only 4 disk drive slots associated with the #6584 or #6594. The other 4 disk slots are always driven by the embedded controller. The highest performance disk controller, the #2757, is expected to offer much more than is needed for just 4 disk drives. It has not been tested in the CEC and is therefore is not supported inside the 520 CEC.

Contact Information

For corrections to this document, contact Jim Cioffi at jacioffi@us.ibm.com . For other technical support or technical sales support, contact your appropriate IBM support channel.