128MB vs 256MB RAID Controllers

Introduction

ICP vortex Computersysteme has recently released the ARC-based ICP9014RO (Single-Channel) and the ICP9024RO (Dual-Channel) RAID Controllers. These are the first of the new generation of ICP RAID Controllers to be launched. Both cards come with the option of a battery back-up (BBU) module that can be purchased separately, as well as:

- Full 64-Bit/133MHz PCI-X support, with up to 1056 MB/Sec bandwidth.
- Integrated ROC (RAID on Chip) architecture for improved performance.
- Onboard 256MB / 64-Bit DDR ECC memory for optimized storage utilization.
- New ICP Storage Manager (ISM) software with graphical ICP RAID configuration utility for easy installation, configuration and management.

This document provides an analysis of the performance differences between these new ICP RAID Controllers with 256MB cache, and the same controllers with only 128MB, and shows why ICP has chosen to offer its controllers with the 256MB cache. The 128MB variants have been specially manufactured in very small numbers, just to enable us to be able to perform a like-for-like comparison with the standard 256MB controllers.

Measuring Performance

The benchmarking programme, Winbench 99 was used to test the performance of these two new ICP products against the 128MB versions. This measures a hard-disk's performance from 2 aspects. The first are low-level measurements, such as CPU utilization, sequential transfer rate, and random access time. Next are the high-level benchmarks, known as Business Disk Winmark 99 and High-end Disk Winmark 99. These attempt to emulate the disk access patterns of typical business applications, by keeping multiple applications open within each suite, and switching tasks between those applications.

The following applications were used in our tests. They were chosen especially because they are very demanding in terms of hard-drive usage:

- Photoshop 4.0 to simulate how a user would move and edit large, high-quality image files.
- Sound Forge 4.0 to show how large audio/video files would be managed.
- Visual C++ 5.0 to replicate the compiling of a large software system.

The individual tasks were run serially, and measured and assigned a score, in addition to a weighted average. These scores were then combined to give us the final results.

WinBench 99 measures the transfer rates of both sequential and random reads and writes. It is these random reads and writes that really slow the system down. Due to the increased head movement of the drives, there are more and more phases where the system needs to wait for the disk to answer. A large cache memory on the controller combined with an optimized algorithm can help to buffer and pre-sort the I/Os to the hard disk(s), thereby substantially optimizing the total throughput.

The Test Configuration

Server	**Pentium 4 Xeon @ 3.6GHz**
System Memory	**512MB**
OS	**Microsoft Windows 2000 Server**
Controller	**ICP9024RO**
Controller Host Interface	**PCI-X (64Bit, 133MHz) 200MHz Mips Processor**
Controller Cache	**128MB*or 256MB (Write Back Enabled)**
Driver Version	**B7367**
RAID Type	**RAID 5 & RAID 0**
RAID Stripe Size	**256KB**
Disk Drives	**Seagate Cheetah 15k.3 Ultra320 SCSI (ST318453LW)**
Disk Capacity (per Drive)	**18.4GB**
Disk Speed	**15,000rpm**
Disk Firmware Version	**0006**
Disk Seek Time	**3.6 m/sec avg.**

* 128MB controller not available. This was manufactured in small number for testing only!

All our testing was carried out using a high-end processor system, so performance was limited by the speed of the drive and the capabilities of the controller, rather than the speed of the system.

In our testing environment, there were 8 drives in each array and for each array, the total capacity was different. In the RAID 0 there were 8 drives at 18.4GB, so the array capacity was 147.2GB. RAID 5 arrays require the capacity of one of the drives for the parity information, so the total array capacity was reduced to 128.8GB in this case.

The stripe size for our testing was done at the firmware's default which is typically 256KB for our PCI RAID controllers. The driver version used for both controllers was B7367.

Our tests were made using RAID 5, as this is the RAID level more likely to be found in a file and application server / database server / web server environment, where there are multiple users and where the amount of memory and cache memory of the controller are even more important.

We also made tests using RAID 0, as this is the RAID level that gives the best performance with data striping, and is most suited to the types of applications we were using, i.e. image editing, video/audio production and editing, pre-press and any applications requiring high bandwidth. RAID 0 is best suited for high performance workstations where redundancy is less important then performance (e.g. after processing, the data is copied back to a server system with some kind of redundancy).

The whole suite of tests was repeated 3 times. The drive is re-partitioned, re-formatted and the system cold-booted between each iteration, which would clear the drive's buffer as well as the OS disk cache.

The Results

Testing was carried out using 3 different configurations:

- 2 ch 8dr5 (NTFS volume)
- 2 ch 8dr0 (NTFS volume)
- 2 ch 8dr0 (FAT32 volume)

	2ch 8dr5 (NTFS volume)			2ch 8dr0 (NTFS volume)			2ch 8dr0 (FAT32 volume)		
	128MB	256MB		128MB	256MB		128MB	256MB	
Winmark 99	20,600	22,900	+11%	20,700	22,500	+9%	25,500	25,400	-1%
Photoshop 4.0	69,500	89,600	+29%	56,400	72,600	+29%	62,300	69,600	+12%
Sound Forge 4.0	193,000	239,000	+24%	210,000	242,000	+15%	241,000	272,000	+13%
Visual C++ 5.0	66,900	82,800	+24%	78,100	88,000	+13%	125,000	132,000	+6%

(Figures are in Thousand Bytes / Sec)

The graphs below help to visualize the differences in these performance figures between the controllers with 128MB and 256MB when tested using Winmark 99 and the other software packages and utilizing the 2ch 8dr5 (NTFS) configuration:

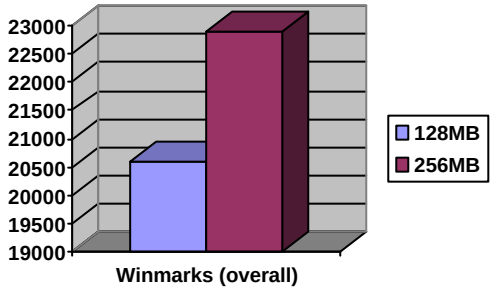

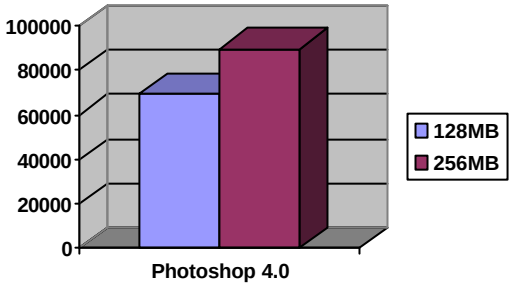

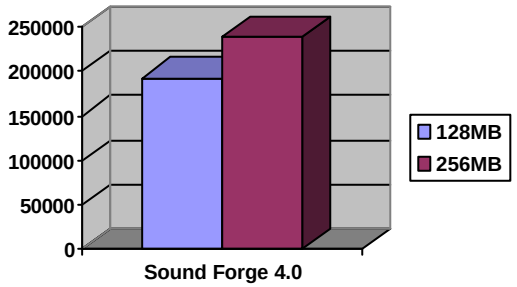

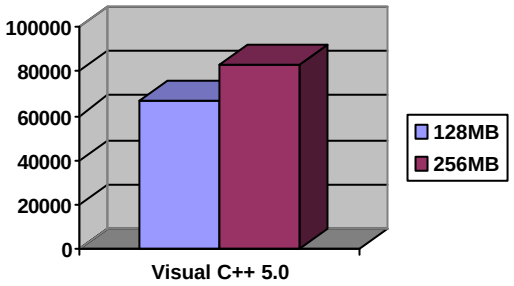

With NTFS, we are seeing a performance gain of ~10% on Winmarks and in some cases, up to 30% on the business applications when comparing the ICP 256MB RAID controller to the 128MB version.

The NTFS file system shows a much better performance gain over all applications, when compared with the tests made using the FAT32 volume.

Even though this benchmark simulates a single user environment, most of the results can be transferred to multi-user environments, where the number of I/Os is higher, and so the performance benefits of a larger cache on the controller are even greater. As discussed earlier, this is especially important when we talk about random I/Os.

The Conclusion

In summary, if you are working with applications which are heavy on hard-drive usage, and so slow down your system, such as those used to compile huge software programs, or handle large audio/video/image files, the amount of cache memory is extremely important to keep the PC running at a reasonable speed. If we translate this example into a server-based environment, where there are multiple users, then the need for increased cache memory becomes even more important. So it follows that you can expect the benefits to be even higher.

In simple terms, the larger the cache memory, the better the performance of the disk array. That's why ICP vortex Computersysteme has chosen to include the 256MB cache memory chip on all upcoming ARC-based SATA II and SAS RAID controllers to be launched in the future.

More product information can be found by visiting www.icp-vortex.com