Partner

RAID 5 for ISS45 — The 4th "R"

MB-ISS45-1142
March 9, 1999

ICL is pleased to announce RAID 5 compatibility for **ISS*45* V*8***. ICL's position is that use of RAID technology will have an enormous benefit to retailers, and that it represents one of the greatest values we can offer.

This bulletin discusses RAID and how it applies and benefits users of **ISS*45* V*8*** systems.

ISS45 VERSION 8 CONFIGURATION

ISS*45* *V8*, like all **ISS*45*** configurations, uses Controller Service replication. In Version 7, there are two essentially identical Master File Server units ("MFS") operating in a "balanced" configuration — that is, both are always operating and available, and there is no separate takeover mechanism required if there is a software, hardware or power failure on one of the units.

In **V*8***, two Controller Services are also used. You could call them an "unbalanced" configuration, however, since only the MFS1 contains the master SQL database tables. MFS2 contains all the "internal" tables used for communication within the **ISS*45*** On Line Transaction Processing System, and all the critical data is therefore on MFS2. The takeover non-mechanism is the same –- that is, if the **V*8*** MFS1 were to leave the configuration, MFS2 would automatically handle all the Controller Service functions. Except …

Except that many of the primary application interfaces may be directly to the master SQL tables (for example via Active-X or ODBC) and these SQL tables may contain additional management information, fields and data that are not in the smaller QDX files that hold data and move it around the store. Many reports will also use the SQL files as their information source, so reporting from the MFS2 will not be in the same format. If you have no MFS1, then your operation is limited to the extent that the master SQL files. You'll be able to run your store, but not at the same level.

So, it was important to find a solution to reinforce the level of Replication for the **V*8*** system.

What about "redundant" capabilities in SQL itself? As you may know, Microsoft SQL (either version 6.4 or 7) running on NT has built-in redundancy features to copy data from one machine to another if they are set up a certain way. Unfortunately, this redundancy capability is useful only in lighter weight applications, and does not provide the kind of bulletproof certainty that retailers have come to expect in redundant or fail-safe systems. ICL has carefully tested SQL redundancy in Wake Forest for **GLOBAL*store*** purposes, and the **ISS*45*** team has also attempted to implement them within **ISS*45* V*8*** without success. These features may be useful in some applications, but not for the constant database pounding typical of multi-user OLTP systems such as PoS.

RAID TO THE RESCUE

RAID stands for Redundant Array of Independent Disks. The original, and brilliant, idea was to group relatively inexpensive drives as a single logical drive to achieve fault-tolerant data redundancy — that is, the ability to recover data that would otherwise be lost by a failed drive —for less than the price of a single larger hard drive.

The basic components of a RAID system include two or more hard drives and RAID software system. Most RAID systems also include one or more cards or array controllers that improve on the performance of the built-in hard disk controller on the PC. RAID 3 and 5 server systems can incorporate one or more hot-swap features — such as removable-drive trays — that allow you to replace a failed drive (and even reconstruct the data on the new replacement drive) without shutting down the RAID system.

Be aware that "software" RAIDs (such as those supported by an NT internal utility) may have some benefits, but performance is ***not*** among them. ICL benchmark tests with the "software" RAID show a substantial decrease in performance and as such these configurations are not recommended for use with **ISS*45*** under any circumstances.

Instead, **ISS*45* V*8*** supports a RAID 5 system that brings remarkable performance, capacity and safety to the PoS system. The recommended **ISS*45*** RAID system contains a high-performance server (see Marketing Bulletin #1047) and a RAID controller managing a three or four-disk array. The RAID system is usually set up as follows:

- One disk is designated as the "System" drive and contains program material only.
- The other three disk drives are configured as the "Data" drives, supported under the RAID controller.
- Data is written to the RAID controller, which splits up the data elements into three smaller pieces:

1. Half of each data element is written to Data Drive 1,
2. The other half of each data element is written to Data Drive 2,
3. A "parity" element is written to the Parity Drive, Data Drive 3, containing the delta between Data Drive 1 and Data Drive 2. The intelligent RAID controller, with either of the pieces of the element from Data 1 or Data 2, plus the Parity element, will be able to reconstruct the full data element.

You'll notice here that no one disk drive can provide the data. Each disk drive has only a portion of the data needed to comprise a complete data element.

- When data retrieval is required, the RAID controller simply grabs the first half of the data from Data Drive 1, the second half from Data Drive 2, assembles them, and provides the reconstituted data element to the CPU as requested. In normal circumstances, the Parity Drive (Drive 3) isn't used or needed for retrieval.

REPLICATION, RESILIENCE AND RECOVERY WITH RAID

The real genius of RAID 5 is when a hardware failure is experienced. RAID is truly "The Fourth R" in an OLTP environment.

- If Data Drive 1 fails, the RAID controller will be able to continue without pause, data or performance loss by simply using the data on Drive 2 plus the Parity drive. If Data 2 fails, the RAID controller can use Data 1 plus Parity to supply the full data. If the Parity Drive fails, then Data 1 plus Data 2 contain all the necessary data.
- The second Big R is Resilience. RAID does all failure handling with no loss of performance, no reconfiguration, no shut down and restart -- nothing. The system gets a failure message about the bad drive to be fixed, that's all.
- And then you get the third Big R - Recovery. With "Hot Swapping," the Recovery is simplicity itself. The RAID controller will not mind if you remove the faulty drive during operations — you can literally open the cabinet without shutting down the processor and pull out the bad drive and replace it. Absolutely amazing. No change of system status is required, no temporary cessation of operations. The RAID machine will continue to operate as if nothing were happening.
- Furthermore, When a replacement drive is inserted the RAID controller will dynamically configure the new or repaired drive and automatically rebuild the data on it to match the current data on the other two drives and return the system to full data replication.

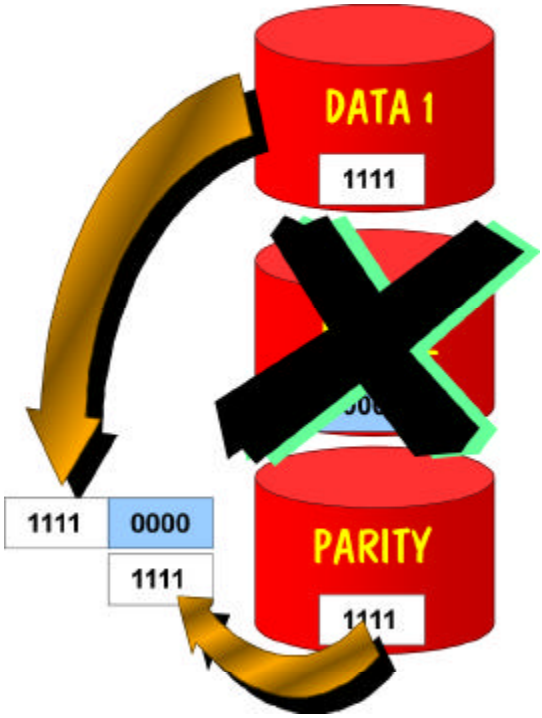

OTHER BENEFITS OF RAID SERVERS

The benefits of RAID 5 technology go well beyond the elegant addition to system security.

- There are other reliability benefits. The system disk for example, will be left for only program tasks which would be far less wear and tear on the mechanism than if it were subjected to the constant database pounding required by SQL system. This alone will normally increase the MTBF and reliability of the system drive.
- You'll notice that since Data Drives 1, 2 and 3 only use half the space each to store their half of the data element, the drive has actually doubled in "capacity" because two drives are used to store one element. For example, if you have an array of 5GB drives, the system will "see" a single monolithic 10GB drive. Again, this is because each 5 GB drive has to store only half the data it otherwise would, so its effective size is doubled. Note that in many cases, when very large drives are required, the multiple drives and controller may cost little more, or maybe even less, than a single very large drive of twice the size.
- Besides the disk size benefit, there are major performance benefits. For example, with **ISS*45* V*8***, careful benchmarks with a 30-terminal test system showed that the RAID system was three times faster than the identical server without RAID. In standard processing of transactions to capacity, the standard server with **ISS*45* V*8*** handled about 200 items per second through the internal and SQL transaction processing system (a stunning number by itself, since that means that even with checkers who can scan at three times the speed of the average checker, **ISS*45* V*8*** could handle the capacity requirements of a 200-lane store). But with RAID, capacity tripled to ***600*** items per second.

- Of course, in practical terms, either benchmark is overkill for handling standard scanning and processing in any normal store. But consider end-of-day, reporting and the like — the blinding speed of the RAID system will enable virtually instantaneous retrieval of reports, even those requiring enormous amounts of file-combing and number crunching. End-of-day can complete in a fraction of the time; since the disk retrieval is so much faster and the data handling is offloaded from the main CPU to the intelligent RAID controller, the main CPU is left free to handle computation and operation while the RAID controller feeds it data even faster.

SERVERS VERSUS PCs

- The ICL RAID system would normally be paired with a similar server, but without RAID for the MFS2 machine. It's important to understand the difference between standard PCs and Server-Class machines. Servers are like MIL-SPEC class machines, designed for much heavier duty and data throughput. Components are matched with closer timing tolerances, and there aren't the little unexplained hardware glitches that you get when you send a lighter weight machine to do the heavy lifting. **ISS*45*** has done remarkably well with PC-Class machines, and the relatively simple V7 architecture and the elegant balanced replication has protected users. But on the heavy SQL and NT platforms the market is demanding, Server class hardware has become a must.
- Modern Servers are designed specifically for NT operation and compliance. Due to this investment, machine-specific oddities are removed, and NT runs cleaner and smoother.
- Servers are turned over in the product line far less frequently, and can provide much better consistency throughout even an extended rollout cycle for maintenance and operations.
- ICL Servers can be outfitted with an optional second Pentium II processor for even faster CPU response.
- RAID 5 adds intelligent hardware to the box, and takes big pressure off the motherboard to perform the disk control. So the system again runs smoother and faster.
- ICL's RAID controller includes a backup power system to back up the cache in the RAID controller — this makes the data writing corruption-proof regardless of power glitches.

STORE CONFIGURATION BENEFITS

As you can see, there are plenty of reasons to use RAID 5 with **ISS*45*** as a true server for the in-store environment. In fact, with the tremendous headroom and bandwidth of a Server-Class machine, use of the server for other in-store applications and electronic payments is an easy choice. So long as the additional applications are "well-behaved" under NT, the entire store will be made more reliable. Here are some additional points that users should consider.

- Spread the RAID around: office applications and especially Electronic Payments are "mission critical" to customer relationships and the operation of the store. For example, in today's retail environment, having never-fail PoS is nice, but not really useful if you can't take ATM or credit cards and your check verification is off-line.
- If three separate PCs are used for PoS, EPS and a Workstation/ISP, there are more than three times the opportunity for a hardware failure than if there is one server unit running the entire operation. With RAID's capability to withstand disk failures, and the Server's

better construction and reliability, a single RAID box will provide far better availability than three separate PCs for these functions.

- Engineers and system integrators and consultants agree that a system with a single high-quality Server is a more robust configuration than a string of networked PCs. Configuration is simpler and with fewer interface layer components, and overall smoothness and data transfer is far superior.
- Cost factors also favor RAID as a superior value. At typical selling prices, a standard RAID MFS1 Server and a matched non-RAID Server for MFS2 will cost in the neighborhood of $6,500 - $7,000 per store more than a pair of high-performance desktops. Given RAID's tremendous benefits of data protection, this is a very inexpensive insurance policy over the life of the hardware platforms. These are not apples-to-apples comparisons, either: the RAID Servers will come with four times the RAM, four to six times the disk space and three times the performance of the best desktops.
- The cost difference is even less or in fact vanishes altogether if — as recommended — the RAID server is used to run other NT functions in the store. This eliminates the purchase requirements for two or three additional PCs and operating systems, client licenses, etc., besides simplifying operation, installation and control.
- In large stores, where the system load is heaviest (greatest performance benefit) and sales are highest (greatest reliability benefit) the cost for RAID will add zero to 3% to the acquisition cost — a no-brainer. Even in the worst case, where use of RAID does not eliminate the extra PCs as described above, specifying the RAID and Server systems will add only 5% or less to the system cost in a typical store.

With this kind of economic profile, it's clear that the **ISS*45*** user stands to gain much more in value than the relatively small cost increment to move to RAID. In addition, **V*8*** customers considering RAID should be aware that **ISS*45***'s largest chain users specify this configuration and that RAID has become our standard validated platform for **ISS*45*** **V*8*** software.

To your success,

Tony

Tony van Seventer
Director: Supermarket Systems