Intel® Storage System SSR212CC

RAID Sets & Volume Capacity – Application Note

Revision 1.0

Server Products Group Technical Marketing

Revision History

Date	Revision Number	Modifications
October 25, 2006	1.0	Release Copy.

Disclaimers

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Storage System SSR212CC may contain design defects or errors known as errata that may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel system boards contain a number of high-density VLSI and power delivery components that need adequate airflow to cool. Intel's own chassis are designed and tested to meet the intended thermal requirements of these components when the fully integrated system is used. It is the responsibility of the system integrator that chooses not to use Intel developed system building blocks to consult vendor datasheets and operating parameters to determine the amount of air flow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible if components fail or the system does not operate correctly when used outside any of their published operating or non-operating limits.

Intel, Pentium, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation.

*Other brands and names may be claimed as the property of others.

Table of Contents

1. The 2 TB limitation

All modern drives are typically organized into blocks of 512 bytes. In the 1980's the SCSI standards team developed an architecture called the Command Descriptor Block, or CDB, which was 10 bytes long. The CDB embedded a block number field that was 4 bytes long. This 4 byte block number represents 4,294,967,296 individual blocks, which represents a total of 2 TB when those blocks are 512 bytes each. Consequently, the maximum raw capacity of a logical drive (or volume) supported on the Intel® RAID Controller SRCS28X is 2 TB.

Although individual volume/logical drive size is limited to 2 TB, larger volumes can created by using disk virtualization utilities of the host operating system, as discussed in this document.

2. Supported RAID levels

RAID levels that are supported on the Intel® Storage System SSR212CC are defined below.

Level 0: Often called "striping," is a performance-oriented data mapping technique. Data written to the array are divided into stripes and written across the disks of the array. This procedure enables high I/O performance at a low cost but provides no redundancy.

Level 1: Often called "mirroring," provides redundancy by writing identical data to each member disk of the array. Level 1 operates with two disks that may use parallel access for high data-transfer rates when reading, but more commonly operates independently to provide high I/O transaction rates. Level 1 provides very good data reliability and improves performance for read-intensive applications but at a relatively high cost. Minimum number of drives is two.

Level 5: RAID Level 5 is the most common type of RAID. By distributing parity across some or all of an array's member disk drives, RAID level 5 eliminates the write bottleneck. The result is asymmetrical performance, with reads substantially outperforming writes. Level 5 is often used with write-back caching to reduce the asymmetry. Because parity data must be skipped on each drive during reads, however, the performance for reads tends to be considerably. Minimum number of drives is three.

Level 10: RAID level 10 employs the features of levels 1 and 0. The advantages are faster data access (like RAID 0), and single-drive fault tolerance (like RAID 1). RAID 10 requires twice the number of disks (like RAID 1), but it offers some performance improvements by striping, then mirroring the striped array. RAID 10 stripes the blocks of data to each RAID 1 array. Each RAID 1 array then duplicates its data to its other drive. Minimum number of drives is four.

Level 50: RAID level 50 employs the features of levels 5 and 0. RAID 50 includes both parity and disk striping across multiple drive groups. RAID 50 is best implemented on two RAID 5 disk arrays with data striped across both arrays. RAID 50 can sustain one to four drive failures while maintaining data integrity if each failed disk is in a different RAID 5 array. Minimum number of drives is six.

As mentioned in section 1 of this document, each Intel® RAID Controller SRCS28X can support multiple volumes, each up to 2 TB maximum, with limitations depending upon hard disk drive capacities, as shown below.

Single Volume Drive Capability

RAID 0 SET SIZE *(per SRCS28X)*	*250 GB*	*400 GB*	*500 GB*	*750 GB*
1	**250 GB**	**400 GB**	**500 GB**	**750 GB**
2	**500 GB**	**800 GB**	**1000 GB**	**1500 GB**
3	**750 GB**	**1200 GB**	**1500 GB**	**2250 GB**
4	**1000 GB**	**1600 GB**	**2000 GB**	**3000 GB**
5	**1250 GB**	**2000 GB**	**2500 GB**	**3750 GB**
6	**1500 GB**	**2400 GB**	**3000 GB**	**4500 GB**

RAID 0 Array Capacity: (Size of Smallest Drive * Number of Drives).

Single Volume Drive Capability

RAID 1 SET SIZE (per SRCS28X)	*250 GB*	*400 GB*	*500 GB*	*750 GB*
2	**250 GB**	**400 GB**	**500 GB**	**750 GB**

RAID 1 Array Capacity: (Size of Smallest Drive) * (Number of Drives) / 2

Single Volume Drive Capability

RAID 10 SET SIZE (per SRCS28X)	*250 GB*	*400 GB*	*500 GB*	*750 GB*
4	**500 GB**	**800 GB**	**1000 GB**	**1500 GB**
6	**750 GB**	**1200 GB**	**1500 GB**	**2250 GB**

RAID 10 Array Capacity: (Size of Smallest Drive) * (Number of Drives) / 2

Single Volume Drive Capability

RAID 5 SET SIZE (per SRCS28X)	*250 GB*	*400 GB*	*500 GB*	*750 GB*
3	**500 GB**	**800 GB**	**1000 GB**	**1500 GB**
4	**750 GB**	**1200 GB**	**1500 GB**	**2250 GB**
5	**1000 GB**	**1600 GB**	**2000 GB**	**3000 GB**
6	**1250 GB**	**2000 GB**	**2500 GB**	**3750 GB**

RAID 5 Array Capacity: (Size of Smallest Drive) * (Number of Drives - 1).

Single Volume Drive Capability

RAID 50 SET SIZE (per SRCS28X)	*250 GB*	*400 GB*	*500 GB*	*750 GB*
6	**1000 GB**	**1600 GB**	**2000 GB**	**3000 GB**

RAID 50 Array Capacity: (Size of Smallest Drive) * (Number of Drives In Each RAID 5 Set - 1) * (Number of RAID 5 Sets). **NOTE: In this 6 drive RAID 50 example, you must first build two 3 drive RAID 5 arrays.**

* RED indicates the 2 TB maximum single volume limit is exceeded, and therefore that drive size should not be configured/supported as a single volume.

In order to span logical volumes (if RAID is accomplished via the Intel® RAID Controller SRCS28X hardware) into a single volume larger than 2 TB, your Operating System must be able to support greater than 2TB logical volumes.

In Windows* Server 2003 SP1 (or later) operating systems, this is accomplished via the Windows Disk Management Console. Linux* operating systems use the Logical Volume Manager (LVM). Please contact your OS vendor for documentation detailing step by step instructions.

3. Best practices for creating Volumes

Creating excessively large single volume file systems in any operating system is a feature that is nice to have, but not really very practical for several reasons: 1) Maintenance tasks are going to take much longer. 2) Any file system issue then effects much more data. 3) RAID re-building/re-stripping efforts can take days, or even weeks to complete (depending on RAID size).

As a *best practice*, allocating multiple smaller volumes of data is a much more conservative and practical approach to managing today's large amounts critical data.

For example, if you want to use all the physical drive space of a 500 GB or 750 GB drive (or larger) in a RAID set, you achieve this by creating multiple 2 TB logical drives/volumes. The equation is:

[Array Capacity£ / 2 TB] = min # of 2 TB volumes

See examples & Table below.

EXAMPLE 1

Example calculation using RAID 5 and six 500 GB disks connected to one SRCS28X , using the equation above (and Table notes below) yields:

[(500 GB) * (6 -1)] / 2 TB max volume size = min # of 2 TB volumes.

[2.5 TB] / 2 TB max volume size = 1.25.

Therefore, the minimum number of 2 TB volumes equals 2, as shown in the table below.

RAID Level	**Disk Drive Capacity**	**# of Disk Drives** (per SRCSX28)	**Minimum # of 2 TB Volumes** (per SRCSX28)
5	500 GB	6	2

NOTE: £

RAID 0 Array Capacity (per SRCS28X): (Size of Smallest Drive * Number of Drives).

RAID 1 Array Capacity (per SRCS28X): (Size of Smallest Drive) * (Number of Drives) / 2.

RAID 5 Array Capacity (per SRCS28X): (Size of Smallest Drive) * (Number of Drives - 1).

RAID 10 Array Capacity(per SRCS28X): (Size of Smallest Drive) * (Number of Drives) / 2.

RAID 50 Array Capacity(per SRCS28X): (Size of Smallest Drive) * (Number of Drives In Each RAID 5 Set - 1) * (Number of RAID 5 Sets).

EXAMPLE 2

Customer A is using twelve 500 GB drives (6 TB of raw storage).

They require redundancy and multiple volumes.

Solution –

- Create 1 TB RAID 5 Volume for company database.
- Create 2 TB RAID 5 Volume for company email.
- Create 2 TB RAID 5 Volume for company documents.

NOTE: 1 TB of storage is consumed with RAID 5 redundancy, per NOTE £ *above*.

EXAMPLE 3

Customer B is using twelve 500 GB drives (6 TB of raw storage).

They require a single 6 TB volume and no redundancy.

Solution –

- Create 2 TB RAID 0 Volume on 1st SRCS28X.
- Create 1 TB RAID 0 Volume on 1st SRCS28X.
- Create 2 TB RAID 0 Volume on 2nd SRCS28X.
- Create 1 TB RAID 0 Volume on 2nd SRCS28X.
- Use your Host Operating System Disk Virtualization Utilities to span all the volumes together.

4. Summary

You can't create a single logical drive (or volume) with more than 2 TB capacity for each SRCS28X, regardless of the type, size, or number of Hard Disk Drives used. However, multiple 2 TB logical drives can be configured, and represented via Operating System Disk Virtualization support as a single volume, as shown in Figure 1 below.

Figure 1.