

ATI SB600 SATA/RAID Features

REVISION HISTORY	3
OVERVIEW	4
SATA FEATURES:	4
JBOD:	4
R0 (STRIPING):	4
R1 (MIRRORING):	4
R10 (STRIPING OVER MIRRORED SETS):	4
ATAPI DRIVES:	4
HOT-SPARES:	4
AUTO REBUILD:	4
HOT-SWAP REBUILD:	4
SMART (SELF MONITORING, ANALYSIS AND REPORTING TECHNOLOGY):	4
TURN ON/OFF DRIVE CACHE FROM THE APPLICATION:	4
SUSPEND/RESUME OF ARRAY MANAGEMENT TASKS:	5
SCHEDULING ARRAY MANAGEMENT TASKS:	5
DRIVE ROAMING:	5
LOAD BALANCING:	5
GREATER THAN 2 TERA BYTE ADDRESSING:	5
RAID LEVEL MIGRATION AND CAPACITY EXPANSION:	5
MULTIPLE LOGICAL DRIVES PER DISK:	5
GIGABYTE ROUNDING:	5
MANAGEMENT SOFTWARE:	6
OPTION ROM BIOS:	6
FACTORY UTILITIES:	6
APPENDIX A	7
PORT MULTIPLIER FUNCTIONALITY	7
NAMING CONVENTION	7
PORT MULTIPLIER CONFIGURATIONS	8
APPENDIX B	9

Revision History

	Description	Updates By	Date
1	First revision	Manjunath Narayanaswamy	09/03/2004
2	Version 1.1 – Added Appendix A and updated to include support of port multiplier and ATAPI devices.	Manjunath Narayanaswamy	01/17/2006
3	Version 1.2 – Updated list of languages supported in the first release. Reduced the list to core language set which is widely expected in the field.	Manjunath Narayanaswamy	07/27/2006
4	Version 1.3 - Added clarification to JBOD definition per AE request	Manjunath Narayanaswamy	08/02/2006
5	Version 1.4 – Correct abbreviation for Chinese Simplified and Chinese Tradition	Lilian Nguyen	01/04/2007

Overview

This document describes and lists all the SATA and RAID features offered by ATI's SB600 south bridge chipsets.

SATA Features:

- Software is AHCI 1.1 compatible and implements Native Command Queuing (NCQ).
- SATA Port Multipliers are supported. Refer to Appendix A for supported configuration.
- GEN 1 and II SATA drives are supported
- Transfer speeds 1.5 and 3 Gbps supported
- Aggressive Link Power Management is supported
- Hot-plug of drives supported

JBOD: Single drive volumes (JBOD concatenation as single volume is not supported).

R0 (Striping): Data is striped across more than one disk to improve performance. This array type does not provide any fault tolerance (i.e. if a disk fails, user may not be able to access logical device anymore). This kind of array is suitable for storing non-critical data and applications where response time is very crucial.

R1 (Mirroring): Data is copied onto two or more disks simultaneously. Most commonly used configuration is 2-way mirroring. This array provides high degree of fault tolerance. Since data is stored on multiple disks, software can do load balancing picking different drives for all READ requests. However, total disk capacity is reduced by half to keep the redundant data. With disk drives getting economical this may not be a significant factor.

R10 (Striping over Mirrored sets): Advantages of R0 and R1 can be combined by striping data across mirrored pairs. This provides fault tolerance and at the same time improving performance on READ requests. This configuration requires minimum of 4 disks.

ATAPI drives: RW CD/DVD drives are supported.

Hot-Spares: One or more disks can be configured as spare disks, which will become active if one of the disks in fault tolerant array fails.

Auto Rebuild: If one of the disks in fault tolerant array fails, then data is recovered from rest of the drives and rebuilt onto a spare member. Once all the data is completely rebuilt, array will return to optimal state providing same level of performance and fault tolerance.

Hot-Swap Rebuild: On a degraded R1 or R10 array, if a good drive is hot-added to missing drive location, REBUILD task will be performed automatically with zero manual intervention.

SMART (Self Monitoring, Analysis and Reporting Technology): Most of the drives in the market have capability to predict drive degradation. RAID engines will query for such cases and prompt user to take a back up and replace the degrading disk to avoid data loss.

Turn on/off drive cache from the application: Some users have a requirement to ensure data is written on the disk once the RAID engines claim WRITE is complete. They can not afford to lose data sitting in drive cache on a power failure. To ensure no data loss they need flexibility to turn off drive cache (at the expense of reduced performance).

Suspend/Resume of Array Management Tasks: Tasks like REBUILD, Synchronization will be resumed from where left off across power cycles and different power modes. With disk drives getting larger in capacity, these tasks might take hours to complete and user might want to shutdown or put the system in Standby mode while the tasks are active.

Scheduling Array Management Tasks: User can schedule REBUILD and Synchronization activities at regular scheduled intervals.

Drive roaming: Recent desktops and workstations come with internal drive bay which can host multiple disks. These drive bays have facility to hot-plug drives and this can result in drive moving to different slots in the bay. RAID engine is tolerant to this kind of roaming and identify members of the array properly. All these tasks shall be performed transparent to user.

Load Balancing: RAID engine will do a load balancing and distribute tasks evenly to members of mirrored pair to improve performance by cutting down response time.

Greater than 2 Tera Byte addressing: The advent of SATA port multipliers and higher disk capacities (400 GB drives in today's market) user's logical volume capacity is expected to exceed 2 Tera bytes in coming years. Currently some of the Linux flavors do have support for addressing higher LBAs and Microsoft will be supporting the feature in Longhorn. ATI offers this feature on Operating Systems which support 2 tera byte volumes.

RAID Level Migration and Capacity Expansion: This feature allows user to morph array to another kind. This can be done online (while user has access to data on the array). This allows user to migrate arrays to add redundancy, capacity, and flexibility for tuning different storage applications.

From	To					
		R0 (1)	R0 (2)	R0 (3)	R0 (4)	R1 (2)
	RAID 0 (1 drive)		Yes			
	RAID 0 (2 drive)			Yes		
	RAID 0 (3 drive)				Yes	
	RAID 1 (2 drive)		Yes			Yes
	RAID 10 (4 drive)				Yes	

Multiple Logical Drives per disk: This feature allows user to configure up to 2 arrays in a given hard disk set. Any remaining capacity after array creation will become independent simple volumes. This is useful especially in the desktop/workstation market where limited disk drives are used to configure arrays both for performance and fault tolerance.

Gigabyte Rounding: The RAID Engine will round individual drive capacities to the nearest common gigabyte value based on a 1000x1000x1000 CHS calculation. This allows for easy interchangeability among various drive vendors. This also provides tolerance to slightly different drive capacities (actual) with same labeled capacity.

Management Software:

- Application has graphical user interface and runs on both Windows and Linux platforms.
- User can use this application to configure, monitor and manage the storage system.
- It is web based and can be used from anywhere on the web.
- Email notification on any disk failure, degradation is supported.
- Localized to 23 different languages – Refer to Appendix B for list of languages.

Option ROM BIOS: Option ROM has a built in configuration tool which can be used to set up bootable arrays, View drive information such as model, capacity, RAID configuration, etc. It is BIOS compliant and can support up to 8 logical drives. Option ROM size is less than 64 KB to suit system manufacturer requirements.

Factory utilities: Linux and DOS based command line scripts are offered to OEM to facilitate production requirements – these are simple applications to create desired RAID configuration.

Appendix A

This section describes functional aspects of ATI SATA/RAID port multiplier solution. This covers attach rate of target devices, system configuration with port multiplier, naming convention etc.

Port Multiplier Functionality

First phase of software release shall support Port Multipliers compliant with AHCI standards version 1.1 and Port Multiplier specification revision 1.1 from Serial ATA committee (<http://www.sata-io.org/>).

Software driver shall implement “**Command based switching**” outlined in the specification and switch to “**FIS based switching**” when it is fully defined in revision 1.2.

Considering memory constraints on the driver, maximum number of devices that will be supported has been limited while making sure it would fit ATI’s target markets.

Software shall recognize and support all Serial ATA devices including S-ATAPI drives. These devices can be directly attached or behind the port multiplier.

All features supported for direct attached devices also applicable for devices behind port multiplier. This includes Hot-plug, Aggressive Link Power Management, and RAID Configuration etc.

Naming Convention

The storage management software, shall list devices in the chronological order as ***Device X***. Device number starts from *Port0*, *Port 1*, *Port 2*, *Port 3*. If Port multiplier is connected to the port, devices behind it shall be enumerated first before moving on. For example, Port multiplier connected to Port 0 shall be enumerated first before attaching device numbers to Port 1.

Port Multiplier Configurations

Configuration		Comments
Number of Port Multipliers	1	One Port Multiplier per system – user should be able to connect it to any of the ports.
Number of devices per Port Multiplier	8	Currently no Port Multiplier in the market supports 8 fan outs. Maximum number of fan outs available is 5.
Number of boot disks	8	Option ROM shall export maximum of 8 boot devices to System BIOS for BBS listing.
Assignment of hard disks among port multipliers		Pick logical devices with lower ID up to 8. For example first list all logical devices connected to Port Multiplier on Port 0, then Port 1 etc.
Max number of logical devices (RAID arrays – virtual disks visible in Operating systems)	8	
Maximum number of physical devices per system	8	Including devices behind port multiplier.

Appendix B

Storage Management Application will be localized to following languages

Abbreviation	Language
CHS	Chinese Simplified
CHT	Chinese Traditional
DEU	German
ENU	English
ESP	Spanish
FRA	French
JPN	Japanese
KOR	Korean
PTB	Portuguese