

SR Performance Updates In Release SR20071022

S. Hopkins
Coraid, Inc.

Summary. Coraid continually improves and enhances the SR Appliance through firmware updates available at the SR support page. Firmware release SR20071022 marks an important milestone in SR development, providing major performance enhancements enabling Coraid's High Performance SRxx61 product series. This release also includes additional features -- most notably, SMART error handling for SR disks. For a complete list of feature changes, please see the SR Changelog at the SR support page¹. This paper compares the performance of the SR152x products using the SR20070924 and SR20071024 firmware releases.

SR Local Rates

The following tables display disk rebuild and parity initialization rates for RAID5, RAID10, and RAID1. RAID10 and RAID1 both rebuild a disk mirror from a single disk; the rebuild effort is essentially the same, but both are presented for clarity.

The following numbers represent the total amount of processed data per second. These rates are sampled using the when command at the beginning of rebuild / initialization and are in units of KB/s. These rates are specifically relevant to the Samsung Spinpoint 500GB hard drives, model SAMSUNG HD512LJ; other disk models may exhibit higher or lower rates based on their capability. As the work proceeds further into the disk(s) the rate will decrease due to the slower disk zones internal to the disk(s). The amount of decrease varies with disk model.

The following tables show the new firmware to have a benefit for all RAID levels, but the greatest benefit is obtained by large RAID5 configurations on the SR1521, where the new firmware provides a 30.4% increase in the rebuild rate and a 19.8% increase in the parity initialization rate.

Release SR20070924:

Model	# Disks in RAID5	RAID5 REBUILD	RAID5 PARITY INIT	RAID10 REBUILD	RAID1 REBUILD
SR1521	14	430,309	406,556	161,917	160,510
SR1520	14	221,118	209,190	158,859	159,184

Release SR20071022:

Model	# Disks in RAID5	RAID5 REBUILD	RAID5 PARITY INIT	RAID10 REBUILD	RAID1 REBUILD
SR1521	14	561,487	487,202	168,764	168,585
SR1520	14	229,926	213,656	168,392	168,827

SR AoE Throughput Rates

The throughput statistics detailed in the following tables were achieved by averaging the results of three independent runs of ddt for the given SR configuration. The system setup and testing approach is as described in the SR Performance Analysis paper available at the SR support page¹.

As shown in the following tables, release SR20071022 outperforms release SR20070924 in most configurations. The greatest improvement is for a 14 disk RAID5 on the SR1521, where the write throughput increases by 37.8%.

1 See Appendix A for URL.

Release SR20070924:

SR1521, two 1GbE links

MTU	KiB / s	RAIDO 15 DISK	RAIDO 4 DISK	RAID10 14 DISK	RAID10 4 DISK	RAID5 14 DISK	RAID5 4 DISK	RAID1 2 DISK	JBOD 1 DISK
9000	WRITE:	186,049.33	117,514.33	110,046.33	65,852.00	157,324.67	83,788.33	34,568.67	37,706.33
	READ:	239,904.00	235,113.00	241,075.00	161,913.00	239,133.67	205,741.67	80,239.67	84,310.67
1500	WRITE:	77,847.67	68,047.67	62,323.67	48,648.33	64,160.33	54,891.00	34,361.33	34,360.00
	READ:	73,386.67	74,711.33	73,295.33	75,113.67	73,138.67	68,574.33	62,600.67	74,965.67

SR1520, one 1GbE link¹

MTU	KiB / s	RAIDO 15 DISK	RAIDO 4 DISK	RAID10 14 DISK	RAID10 4 DISK	RAID5 14 DISK	RAID5 4 DISK	RAID1 2 DISK	JBOD 1 DISK
4200	WRITE:	96,036.67	82,390.00	51,726.67	42,204.00	87,459.33	61,653.33	33,154.00	34,836.33
	READ:	117,887.00	118,373.67	119,161.33	114,348.00	118,711.67	110,657.33	79,436.00	78,973.33
1500	WRITE:	50,069.67	48,702.67	41,735.67	36,773.67	45,179.33	40,088.33	32,056.33	34,218.67
	READ:	54,841.33	55,225.67	54,912.00	55,234.33	54,483.67	52,300.67	47,186.67	56,502.33

Release SR20071022:

SR1521, two 1GbE links

MTU	KiB / s	RAIDO 15 DISK	RAIDO 4 DISK	RAID10 14 DISK	RAID10 4 DISK	RAID5 14 DISK	RAID5 4 DISK	RAID1 2 DISK	JBOD 1 DISK
9000	WRITE:	239,333.33	106,840.33	162,068.67	71,196.67	216,853.67	92,672.33	36,525.00	38,057.67
	READ:	240,682.33	240,891.33	235,480.33	161,025.00	236,099.00	212,946.67	80,943.00	84,789.67
1500	WRITE:	114,288.67	97,436.00	65,229.67	54,341.00	70,224.33	59,805.33	35,432.00	36,611.00
	READ:	112,122.00	110,449.33	95,800.67	92,239.67	98,338.00	91,226.67	80,979.33	84,619.33

SR1520, one 1GbE link¹

MTU	KiB / s	RAIDO 15 DISK	RAIDO 4 DISK	RAID10 14 DISK	RAID10 4 DISK	RAID5 14 DISK	RAID5 4 DISK	RAID1 2 DISK	JBOD 1 DISK
4200	WRITE:	104,253.00	100,832.00	52,251.67	49,750.33	96,317.00	72,052.00	34,966.00	36,765.67
	READ:	119,956.33	120,094.00	119,336.33	117,316.67	119,962.33	117,551.33	77,844.00	82,833.67
1500	WRITE:	71,905.67	67,677.00	41,572.00	38,183.67	42,842.33	40,682.33	33,684.00	36,431.00
	READ:	70,285.33	70,654.00	62,926.67	60,193.33	62,777.33	59,120.33	52,552.00	69,435.00

¹ For an explanation of why the SR1520 performs best with one interface and an MTU of 4200, please see Application Note ANSR001 available at the SR support page (URL available in Appendix A).

Appendix A - References

The SR support page includes the SR firmware, user manual, and related docs:

<http://www.coraid.com/support/sr/>

Please e-mail support@coraid.com with any questions or comments.