

# Overland SnapSAN S2000 Storage for Virtual Servers

## A Technology Validation Report



for  
**Overland Storage**

## About DeepStorage Lab Validation Reports

**D**eepStorage Lab Validation Reports are designed to help IT professionals evaluate emerging products and technologies. DeepStorage Lab Validation Reports are based on our hands-on testing and over 25 years of experience making technology work in the real world.

Our philosophy of real world testing means we configure systems as we expect most customers will use them thereby avoiding “Lab Queen” configurations designed to maximize benchmark performance. It’s frankly too common for published benchmark results to be generated from systems that have been tuned to deliver maximum performance by testing configurations that no customer would use in the real world. These tricks include dedicating large numbers of drives to support small data sets, wasting capacity or using RAID 0 to stripe data across drives without any data protection.

This report was sponsored by Overland Storage. However, all testing for this report was designed and performed at DeepStorage’s Lyndhurst, NJ labs by DeepStorage’s staff. There were no vendor representatives present during testing and DeepStorage always retains final editorial control over our publications.

## Contents

iThe Bottom Line	I
Server Virtualization Changes SME Storage	I
Virtual Server Mixed Workloads Demand Performance	2
Introducing Overland's SnapSAN S2000	3
Key Features:	3
High Availability Features	3
Living with the S2000	4
Windows and Hyper-V Integration	5
vSphere Integration	5
Performance testing	6
IOmeter Testing Results	6
JetStress Testing	9
Mixed Workload Testing	10
Overall impressions and takeaways	12
Test Configuration	13



## Server Virtualization Changes SME Storage

While the Fortune 500 have used storage area network (SAN) technology to share high performance, high reliability storage arrays amongst servers for over a decade, smaller organizations have typically used the simpler direct attached storage (DAS) model.

With today's large disk drives, the DAS model can result in lower disk utilization as free space on one server is "trapped" where other servers can't access it. This is usually offset for smaller organizations by the lower cost of internal RAID controllers and disk drives. Even when external disk shelves, known in the industry as JBODs for 'just a bunch of disks', are needed for servers that outgrow their internal drive capacity the cost per gigabyte is usually lower than for equivalent capacity on a SAN array.

While DAS can be cost effective, it isn't very flexible as DAS disks are available to just one server. As server virtualization technologies, including VMware's vSphere and Microsoft's Hyper-V, make their way into even small and medium enterprise data centers, administrators are discovering that some of the most attractive features of these platforms aren't available without shared storage.

### The Bottom Line

The SnapSAN S2000 by Overland Storage is an easy to use, cost effective, expandable iSCSI storage system designed for the small to medium business market.

- The 2U 12 drive S2000 is expandable to 96 SAS and/or SATA drives through SAS attached expansion chassis.
- The SnapSAN's vCenter plug-in for array management vastly simplifies day-to-day storage management in VMware environments.
- The S2000 provides unexpectedly good performance delivering 5X the performance of internal mirrored drives.
- It also performed well in mixed workload testing supporting the following loads simultaneously from four virtual servers using just 12 SATA drives:

Hypervisor	Heavy Exchange Users	Simulated OLTP IOPS	Simulated web server IOPS	Simulated file server IOPS
vSphere	760	256	156	201
Hyper-V	375	230	153	180

- Snapshot capability is included in the base price with mirroring for high availability and point in time replication for disaster recovery available as options.
- S2000 performance and features are comparable to products from major vendors costing twice as much.

With shared storage, organizations can build central libraries of virtual server templates vastly simplifying the process of deploying new virtual machines. Even more significantly, administrators in shared storage environments can migrate running virtual servers between host systems so the host can be maintained or upgraded – a feature VMware calls vMotion.

VMware has extended this Vmotion functionality to provide high availability automatically migrating virtual servers to a new host when a host goes offline and dynamic resource scheduling which balances the load across a cluster of physical host servers migrating virtual machines from one host to another as workloads shift.

Large organizations have built storage area networks, SANs, using Fibre Channel, a networking technology designed especially for SANs. While Fibre Channel SANs can support huge systems with high performance demands, they require organizations to build an additional network infrastructure in their data centers and learn a new networking technology that seems arcane to administrators used to Ethernet and IP.

By comparison, iSCSI SANs like those that include the SnapSAN S2000, leverage gigabit Ethernet by sending shared storage traffic via IP. Most operating systems and hypervisors, including Windows and vSphere, come with software initiators that use standard Ethernet ports and IP for storage traffic. On the other hand, connecting a server to a Fibre Channel SAN requires a host bus adapter card in the server and expensive Fibre Channel switches.

Once an organization installs a traditional SAN, they soon discover that it is designed to allow storage specialists to fine tune RAID levels and logical disk layouts for particular workloads. This allows the administrator to optimize random I/O performance for data bases or maximum capacity for archival data. It can take many hours to get a traditional array installed and configured to the point where servers can use its capacity. This process takes many mouse clicks, frequently in separate user interfaces, to add a new server to the SAN.

Smaller enterprises can't afford dedicated storage staffs but they still need SAN storage to get the most out of their virtualization investment. For these users, shared storage systems need to be easy to use. In addition, since they will most likely be managed by the server administrators, they should integrate with the server OS management tools.

### **Virtual Server Mixed Workloads Demand Performance**

Virtual server environments also present a performance challenge to storage systems. A physical database or Exchange server might use a RAID 1 volume for its database, which requires random I/O performance and RAID 5 volumes for its operating system and log files that perform mostly reads and sequential I/O. The virtual version of that server will share a data store with other virtual servers with different storage workloads. Since the requests for all these different workloads are intermingled, the storage system never really sees sequential I/O.

## Introducing Overland's SnapSAN S2000

Overland Storage's SnapServer group pioneered NAS for the small business market beginning in 1998 and the S2000's design reflects the many years of experience they've had with these customers. While enterprise storage administrators value performance, reliability and the ability to fine tune their systems through a command line interface, SMB system administrators, who don't configure storage on a daily basis, need systems that are easy to use. The S2000 delivers the ease of use this market segment demands.

### Key Features:

- Quad-core Xeon processor
- 2U rack mount chassis
- 12 hot-swappable bays for SAS or SATA drives
- Support for up to 7 12-bay expansion chassis connected via 4 channel SAS
- 2 Gigabit Ethernet ports
- Up to 8 optional Ethernet ports
- Hardware assisted RAID levels 0, 1, 5, 6, 10, 50, 60 with optional hot spares
- Snapshot support including Windows VSS support
- Synchronous mirroring with automatic failover
- Failover via multipath
- Access control by iSCSI name (iqn) or CHAP
- 3 year warranty

### High Availability Features

Enterprise storage systems are routinely built around custom hardware with features to ensure that they remain operating even when a critical component like a disk controller fails and to protect data from larger disasters that may disable a data center. These features are less common on systems like the S2000 that are based on standard server architectures so they can bring affordable storage to the masses.

The SnapSAN S2000 can be configured to synchronously mirror data for selected logical volumes between a pair of arrays. When data is written to a mirrored volume, the primary array writes to its disks and simultaneously sends the request across a gigabit Ethernet link to its partner appliance for the volume. When both arrays have written the data to disk, the primary system sends an acknowledgement to the server that made the initial write request. Data is therefore kept in perfect synchronization between the two arrays.

## Overland SnapSAN S2000 Storage for Virtual Servers

The server is connected to the volume on both arrays through Windows or VMware's built-in multipath I/O feature that also ensures the server can access the volume if an Ethernet link between the server and the arrays fails. Should the primary array fail, the server will continue to access the volume from the partner array.

In our testing, mirrored S2000s failed over transparently as promised. We started an IOMeter session accessing a mirrored volume and, while it was running, disconnected the Ethernet cable from one of the arrays. The benchmark continued running as if nothing had happened.

Mirroring is a great feature but the fact that it requires low latency gigabit Ethernet connections between the SnapSANs as well as between the servers and both arrays

**S2000 mirroring failed over to the working system transparently in our testing**

limits it to single datacenter or campus environments. To protect data across longer distances, the S2000 can use point-in-time replication. The system journals changes made to a replicated volume on a dedicated journal disk, then periodically sends those changes asynchronously to its partner across whatever IP connection is available.

Using snapshot technology can vastly reduce the amount of data sent across a WAN link between an organization's primary and disaster recovery sites since most sets of data only change 2-3% a day. The problem is if you size your WAN link to match the data change rate, the initial replication can take several days or more. SnapSAN users can use external USB hard drives to seed the replication process – a much simpler process than the alternative.

### Living with the S2000

Overland Storage provided us with a pair of S2000 appliances each equipped with 12 1TB SATA drives and software licenses for both mirroring and replication so we could run them through their paces.

Installing the S2000 was a snap, no pun intended. It acquired an IP address from the lab's DHCP server and the *SnapServer SAN Manager* we installed on our Windows management server discovered them. After configuring all the storage in each test system into a storage pool, and setting a system password, we were mapping targets to windows drive letters in just a few clicks. Those using the S2000 in environments without a DHCP server can set the system's IP address by connecting a monitor and keyboard.

For those organizations using Linux, MAC OS or any of the other non-Windows operating systems, the S2000 includes a web-based user interface which has essentially the same functionality as the Windows version for provisioning and managing storage.

## Windows and Hyper-V Integration

For Windows users *SnapServer SAN Manager* is the primary interface for management of the S2000. Where most iSCSI arrays provide a GUI for management tasks like creating logical volumes and mapping them to servers by iqn, when run from the server accessing the volume, *SnapServer SAN Manager* goes the next step by mapping the newly created logical volumes to Windows drive letters.

This takes several steps out of the usual process of creating a logical drive, configuring access control based on the IP address or iqn of the servers that will be accessing it, and then connecting the Windows server through the Microsoft iSCSI initiator control panel applet.

*SnapServer SAN Manager* also includes a Volume Shadow Copy Services (VSS) provider. VSS coordinates the snapshot process so a backup application, or the S2000's scheduler, can take a snapshot and the applications on the Windows server whose data is being snapped will flush their buffers and create a consistent data image. The S2000 can support up to 256 snapshots per array and unlike many other storage systems, administrators don't have to pre-allocate disk space for snapshots.

## vSphere Integration

While Snap's SAN Manager and/or web SAN Manager are the S2000's primary management interfaces, once the system is configured system administrators can do their day-to-day management through a vCenter plug-in.

After installing the Windows SAN Storage Manager on our vCenter server, we could access the plug-in from any vCenter client. We then created VMFS data stores following a simple six-step wizard in just a few seconds. The wizard walked us through the process of creating a simple data store and made it immediately available to the hosts we selected. When we selected both appliances as the location of a data store, the wizard set up a mirroring relationship between them and verified that our VMware servers had sufficient paths to both arrays to support failover.

**We then created VMFS data stores following a simple six-step wizard in just a few seconds.**

On a conventional system, we would have had to use the array user interface to create logical volumes (LUNs), then switch to vCenter to create a VMware data store and attach each ESX or ESXi server to that data store. Moving array management into a vCenter plug-in like Overland's eliminates many steps in the storage provisioning process that typically include bouncing back and forth between multiple user interfaces.

### Performance testing

While products targeted at the SME market have to be easy to use, an easy to use product that doesn't deliver sufficient performance will cause no end of problems for the system administrators that install it. We therefore ran the S2000 through a gauntlet of benchmarks to see just how well it performed.

In keeping with the DeepStorage.net philosophy of real world testing, we configured the S2000 with a single RAID 10 storage pool. It's frankly too common for published benchmark results to be generated from systems that have been tuned to deliver maximum performance by using configurations that no customer would use in the real world. These tricks include dedicating large numbers of drives to support small data sets wasting capacity or using RAID 0 to stripe data across drives without any data protection.

Readers should also note that our testing was performed on an S2000 with 12 1TB 7200RPM SATA drives which we, and Overland, believe is the most attractive configuration for SMB users. It is also available with 600GB, 15,000RPM drives which reduces the system capacity but should almost double random I/O performance.

### IOmeter Testing Results

Our first round of performance testing used the industry standard IOmeter benchmark. This was originally developed by Intel and is now maintained as an open source project at IOMETER.ORG. A test engineer using IOmeter can perform high levels of I/O against raw disks or file systems. IOmeter profiles allow the engineer to define the data block sizes and the balance of read vs. write and sequential vs. random operations.

While we at DeepStorage.net test with over a dozen IOmeter profiles, many of these are edge cases, like 1 byte random writes or 2MB sequential reads, that are designed to identify potential problems with the system under test or with our test configuration, but do not represent real world conditions.

Over the years, we've found three well known profiles best represent the real world performance of the systems that have come through our lab:

- » The file server profile was developed by Intel based on real world Windows file server access patterns. It includes a mix of block sizes and combined reads, writes, random and sequential IO.
- » The web server profile, also developed by Intel, is dominated by random reads as it emulates a web server retrieving HTML and image files as users load pages.
- » The 4K OLTP profile emulates a database system, like SQL Server 2000 or 2005, that reads and writes 4K pages to its database while running online transaction processing applications like order entry or ERP.

In the market the S2000 is designed to serve, most system administrators will configure it as a single storage pool for reasons of both storage efficiency and simplicity. We ran the IOmeter benchmarks against both an S2000 with a RAID 60 pool, which is our general recommendation as a good balance between storage efficiency and performance,


## Overland SnapSAN S2000 Storage for Virtual Servers

and as a RAID 10 pool which improves random I/O performance, especially on writes, while making less space available for user files.

We also ran the benchmarks against a mirrored pair of 10K RPM enterprise disk drives in the server running Iometer to provide a reference for administrators that are familiar with the performance of local RAID.

Intelligent disk subsystems like the S2000 generally perform better when they can process multiple I/O requests simultaneously. This allows them to pre-fetch data into their cache and process their request queue in the order that will satisfy all requests in the shortest period of time even if it means not processing all requests in the order they were received.

We ran our tests with one, eight and sixteen outstanding IO requests and as expected the S2000's performance was much better with more outstanding IOs. Since the S2000 will usually be installed as a shared storage system accessed by multiple servers, performance with multiple outstanding IOs is a good measure of real world performance.



**Figure 1 - Iometer 4k OLTP IOPS**

Note that the locally mirrored drives deliver about the same performance regardless of the number of outstanding IOs.

### Why this matters:

By emulating a SQL Server running an ERP or other online transaction processing application the 4K OLTP Iometer profile emulates most organizations most performance sensitive, and mission critical application. By wide striping and intelligently managing its SATA drives the S2000 provides enough performance for an SME's most demanding applications at an attractive price.

## Overland SnapSAN S2000 Storage for Virtual Servers

Outstanding IOs	1	8	16
Mirrored 10K Drives	151.5	205.4	214.4
S2000 RAID 60	182.9	743.2	990.9
S2000 RAID 10	158.9	804	1594

**Table 1 – 4K OLTP profile IOPS**

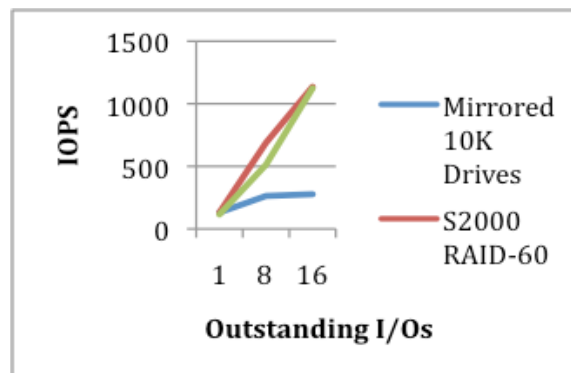
In this most demanding of the profiles, the S2000 delivers 4-5 times the performance of mirrored disks.



**Figure 2 - IOmeter File Server Pattern IOPS**

Outstanding IOs	1	8	16
Mirrored 10K Drives	134	222.6	226
S2000 RAID 60	155	705	1006.9
S2000 RAID 10	131	650.6	1372.5

**Table 2 – IOmeter File Server Pattern IOPS**



**Figure 3 - IOmeter Web Server Pattern IOPS**

	1	8	16
Mirrored 10K Drives	131.4	263.3	277
S2000 RAID 60	131	688.4	1135.6
S2000 RAID 10	115.8	512.9	1124

**Table 3 - IOMeter Web Server Pattern IOPS**

Even when running RAID 6 to balance usable space with data protection, the S2000 delivered at least 4.5 times the performance of the reference mirrored 10K RPM drives. When configured in RAID 10 mode, it delivered over seven times the performance of the mirrored drives with just six times as many slower drives – an admirable performance.

## JetStress Testing

IOMeter allows a test engineer to control the size and distribution of reads and writes to the system under test. However it can only approximate the I/O load of a real application. While the 4K OLTP IOMeter pattern attempts to model database applications like SQL Server or Exchange, real database applications simultaneously perform 4 or 8K page reads and writes to and from the database as they are also writing to their transaction logs sequentially.

### Why this matters:

Many Exchange administrators are convinced that they have to use expensive 10K or 15K RPM drives for Exchange. The S2000 successfully ran JetStress simulating 1100 heavy mail users with 7200RPM SATA drives.

Microsoft's JetStress uses the same Jet database engine as Exchange server to test the ability of storage systems to support Exchange server. Like a production Exchange server, it performs random I/O to its databases and sequential I/O to transaction files. Where IOMeter simply measures the total amount of disk I/O the subsystem can perform, JetStress will fail a test run where database I/O latency exceeds 20ms.

We ran JetStress on an S2000 configured with a RAID 10 data pool supporting two databases. We ramped up the number of heavy Exchange users, defined by Microsoft as .45 I/O operations per second per user, until JetStress reported the system's latency exceeded the 20ms limit at any time over a two hour test run.

The S2000 successfully supported 1100 heavy users in our testing.

## Mixed Workload Testing


In SAN, and especially virtual server, environments it's rare to see a storage system serving an isolated workload like Exchange or SQL Server. Instead, multiple virtual servers each running a different workload all share the storage system. The DeepStorage.net mixed workload benchmark emulates this real world workload.

The DeepStorage.net mixed workload methodology uses four virtual servers one each running:

- » IOmeter's web server profile
- » IOmeter's file server profile
- » IOmeter's 4K OLTP profile
- » JetStress

To determine the aggregate performance of the system under test, we start by running the four benchmark programs each at a level significantly less than the performance the system achieved in individual tests. We then increase the workload by increasing the number of outstanding I/Os in IOmeter and the number of users JetStress is emulating until we reach the point where JetStress reports excessive I/O latency of over 20ms.

We ran this mixed workload benchmark under both Windows 2008 R2 with Hyper-V and VMware's vSphere 4.0 using VMware ESXi. As expected the performance of each workload was reduced from when that workload had exclusive use of the S2000 in both cases the S2000 handled the mixed workload with aplomb.



**Figure 4 - Mixed Workload Benchmark IOPS**

Workload	IOPS under vSphere	IOPS under Hyper-V
4K OLTP	256	230
Web	156	153
File	201	180
JetStress	342 (760 heavy users)	170 (375 heavy users)
Aggregate	955	733

**Table 4 - Mixed Workload Benchmark IOPS**

**Why this matters:**

Our multiple workload testing demonstrates that the S2000 is capable of being the primary storage system for an organization with 500 knowledge workers supporting all of their applications simultaneously.

### Overall impressions and takeaways

With the S2000, Overland Storage demonstrates that the standard server architecture can make a fine storage system when the right components, and most importantly the right software, are skillfully integrated. It delivers a feature set and performance comparable to systems like HP's P4000, powered by the software HP acquired with Lefthand Networks, at roughly half the price.

In our testing, the system performed better than a generic 12 drive SATA array could be expected to reflecting intelligent use of wide striping and cache memory. We would feel comfortable recommending the S2000 to organizations with several hundred knowledge users.

While many systems in this class have limited expandability, by comparison the S2000 can expand to 96 drives through SAS attached JBODs. In addition, the PCI-e architecture allows users to add additional gigabit Ethernet or, as they're ready for it, 10 gigabit Ethernet cards to their systems in the field as their I/O requirements change.

Most significantly, Overland Storage's vCenter plug-in makes the day-to-day management of the SnapSAN S2000 even simpler than managing direct attach storage. Administrators with no storage experience can be creating VMware data stores, even high availability mirrored data stores, in just a few minutes.

While we're fond of the S2000, no product is perfect and we would like Overland Storage to make a few improvements:

- Overland Storage hasn't yet released a VMware Site Recovery Manager module to automatically shift workloads to a disaster recovery site.
- Overland has done a great job with the vCenter management plug-in consolidating many array and vSphere storage management tasks into just a few mouse clicks but building clustered Hyper-V disks isn't as well integrated. SnapServer SAN Manager should be extended to automatically create cluster shared volumes and Hyper-V virtual disks to achieve the same ease of use.
- While the S2000 can be expanded to 96 disk drives, there is only one 4 channel SAS connection that's daisy chained through all the E2000 expansion JBODs. We'd prefer redundant links as an intermediate step before proceeding to mirroring multiple S2000s.
- Support for NIC teaming would improve bandwidth and simplify multi-pathing.

## Test Configuration

Overland Storage provided DeepStorage.net a pair of S2000 arrays each equipped with twelve 1TB SATA disks. Each was also equipped with additional gigabit Ethernet ports.

We connected both arrays to an Extreme x480 gigabit Ethernet switch. The x480 is a 48 port, low latency, enterprise-class switch commonly used for top of rack switching in large datacenters.

IOmeter testing was performed using a Dell PowerEdge 2950 server with dual quad-core Xeon E5310 processors and 16GB of memory running Windows 2008 R2 Standard Edition. The 2950 had a single gigabit connection to the iSCSI network and a gigabit connection to the management network.

VMware management, JetStress and the mixed workload testing were run on a Dell R710 server with Dual Xeon E5620 “Westmere” processors and 24GB of memory running VMware ESXi 4.0U1. Each virtual server on the R710 ran Windows Server 2008 R2 Standard edition with 4 virtual CPUs and 6GB of memory.

Further details including the IOmeter and JetStress configuration files and benchmark output is available from the DeepStorage.net website.

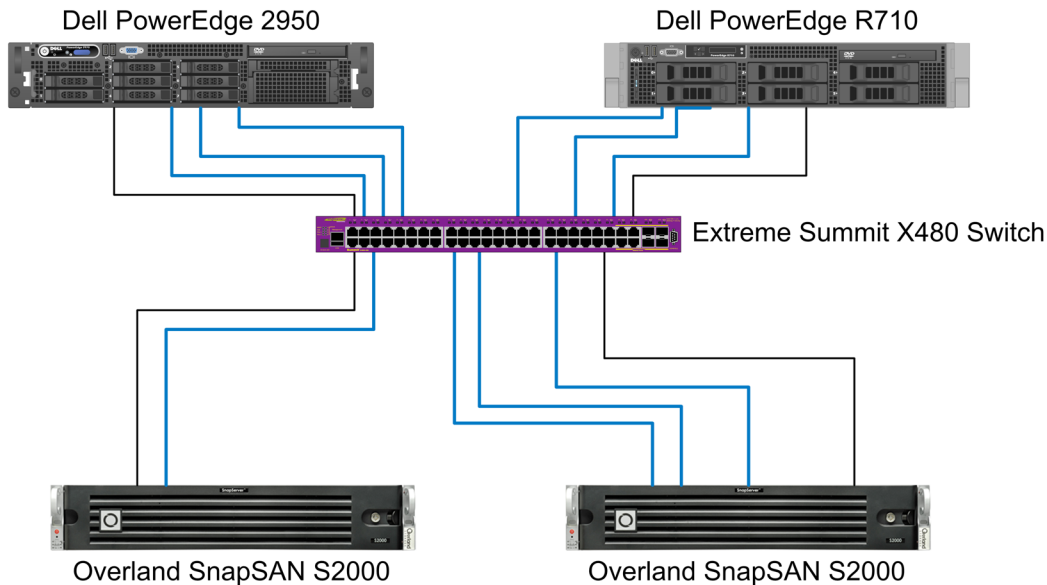


Figure 5 - Test Configuration