www.violin-memory.com
Violin Memory, Inc.
sales@violin-memory.com
+1.888.9VIOLIN x10

VIOLIN 1010 – MEMORY APPLIANCE

FLASH CONFIGURATION AND PERFORMANCE BRIEF

The Violin 1010 is the industry's most scalable and highest performing memory appliance. Memory appliances are designed to eliminate performance bottlenecks that plague I/O-bound applications by providing high performance solid state storage using either DRAM or flash memory. This document explains the various Flash configuration options for the Violin 1010 Memory Appliance and provides information on how it performs in these configurations.

The performance data included in this document is associated with release 2.1 of the Violin 1010. This release commenced customer trials in December 2008. Further performance improvements are expected in later releases. Please contact for Violin for the latest information.

1. CAPACITY

The Violin 1010 is configured with up to 62 Flash Violin Intelligent Memory Modules (VIMMs); VIMMs support 32GB or 64GB of SLC (Single Level Cell) flash memory. The total capacity of the system is calculated as follows:

Raw Capacity = Number of VIMMs x Size of VIMMs

Raw capacity for the Violin 1010 is expressed as traditional storage products (see www.snia.org) are marketed using powers of ten; Billions or Trillions (e.g. 1TB = 10^9 bytes of storage).

Confusingly, flash memory is marketed in the traditional DRAM model using powers of 2 (e.g. 1TB = 2^{40}). For instance, the 64GB VIMM is 2^{36} bytes or 68.7GB of "storage measured" capacity.

A table of Violin 1010 Flash configurations and their raw capacities is below:

Violin 1010 Configuration	VIMM Size (GB)	VIMM Number	Raw capacity	
			GB (2^9)	Trillion Bytes (10^{12})
V-1010-320G-NAND	32	10	320	0.34
V-1010-672G-NAND	32	21	672	0.72
V-1010-1.3T-NAND	32	42	1344	1.44
V-1010-2.0T-NAND	32	62	1984	2.13
V-1010-2.6T-NAND	64	42	2688	2.89
V-1010-4.0T-NAND	64	62	3968	4.26

Table 1: Raw Capacity of the Violin 1010

1.1. Useable Capacity

The raw capacity of any enterprise flash system is not entirely available for user data. Enterprise flash systems reserve some of their capacity to improve the reliability, lifespan and performance of the system. The following overheads reduce the useable capacity:

RAID Overhead: The Violin 1010 uses the Violin Flash RAID algorithm that both (1) provides protection against block, die and module failures; and also (2) eliminates many of the delays associated with flash memory, particularly the erase times. This algorithm uses a RAID stripe of 4 VIMMs with 1 parity VIMM, which means that 20% of VIMMs are not available for user data. This is considerably more efficient than most other flash systems, which use RAID 1 in which 50% of the modules are not available for user data.

Spare Modules: In the rare event that a flash module fails, the Violin 1010 allocates spare modules to transparently rebuild the RAID groups without the applications shutting down or losing any data. Typically, 5% of the VIMMs are spares. Other systems have either no spares (these systems cannot operate safely after a failure) or allocate over 10% of the modules to be used as spares.

Performance Overhead: Enterprise-grade flash storage systems use a page mapping algorithm to provide better Write IOPS (Input/output Operations Per Second). This algorithm requires a "garbage collection" process which consumes some of the flash capacity, referred to as the performance overhead. The Violin system allows the performance overhead to be configurable within a range between 9% and 50%. Configured with 10% performance overhead, the Violin 1010 performs similarly to other flash systems with 50% performance overhead because of the massive parallelism of the Violin Switched Memory (VXM) architecture and the Violin Flash RAID algorithm.

The table below shows the maximum useable capacity of the Violin 1010. The useable capacity decreases and the performance increases when the performance overhead is increased.

Violin 1010 Configuration	VIMM Size (GB)	VIMM Number	VIMM Spares	VIMM RAID Overhead	Performance Overhead (Minimum)	Useable capacity	
						GB (2^9)	Trillion Bytes (10^{12})
V-1010-320G-NAND	32	10	0	2	9%	233	0.25
V-1010-672G-NAND	32	21	1	4	9%	466	0.50
V-1010-1.3T-NAND	32	42	2	8	9%	932	1.00
V-1010-2.0T-NAND	32	62	2	12	9%	1398	1.50
V-1010-2.6T-NAND	64	42	2	8	9%	1864	2.00
V-1010-4.0T-NAND	64	62	2	12	9%	2796	3.00

Table 2: Maximum Useable Capacity of the Violin 1010

2. VIOLIN 1010 PERFORMANCE

The configuration of the Violin 1010 can be varied to find the appropriate trade-off between performance and useable capacity. Performance varies based on the capacity of the system (more VIMMs means more performance) and the performance overhead reserved for page mapping and garbage collection.

Garbage collection is the process of reclaiming pages in flash pages that are no longer used for active data. Garbage collection requires large blocks of flash to be read, the active data pages written to a new block and the inactive pages are effectively recovered for new writes. The process requires spare capacity to be available within the system. With a greater the amount of spare capacity (performance overhead), the garbage collection process is more efficient and the sustained Write IOPS are higher.

All of the performance data provided on the Flash configuration options for the Violin 1010 Memory Appliance assumes 4KByte Block sizes and random access patterns. The system is engineered to support larger blocks as N x 4K blocks.

When the Violin 1010 is configured with 34% performance overhead, very high performance is achieved with over 200K sustained (all day) random Write IOPS. In contrast, Enterprise SSDs and PCIe cards support less than 10K sustained random Write IOPS. The table below provides the key performance metrics for this configuration:

Violin 1010 Configuration	**VIMM**	**VIMM**	**Useable capacity 34% Perf. Overhead**		**Random or Sequential Performance (4K IOPS)**		
	Size (GB)	**Number**	GB (2^9)	Trillion Bytes (10^{12})	Read	Peak Write	Sustained Write
V-1010-320G-NAND	32	10	168	0.18	200K	100K	65K
V-1010-672G-NAND	32	21	336	0.36	345K	200K	115K
V-1010-1.3T-NAND	32	42	672	0.72	345K	219K	204K
V-1010-2.0T-NAND	32	62	1008	1.08	345K	219K	219K
V-1010-2.6T-NAND	64	42	1344	1.44	345K	219K	219K
V-1010-4.0T-NAND	64	62	2016	2.16	345K	219K	219K

Table 3: High Performance Configurations of the Violin 1010

The data in this document assumes a Violin 1010 is connected via a PCIe x8 interface. If connected via a PCIe x4 interface, limits are imposed by the PCIe interface and protocol. For reads this limit is 215K IOPS and for writes this limit is 145K IOPS.

The following graph shows the performance of the Violin 1010 with PCIe x8 interface, 40 VIMMs and a high performance overhead (34%). The bandwidth for all block sizes above 4K is relatively

stable as confirmed in a customer trial of the Violin 1010.

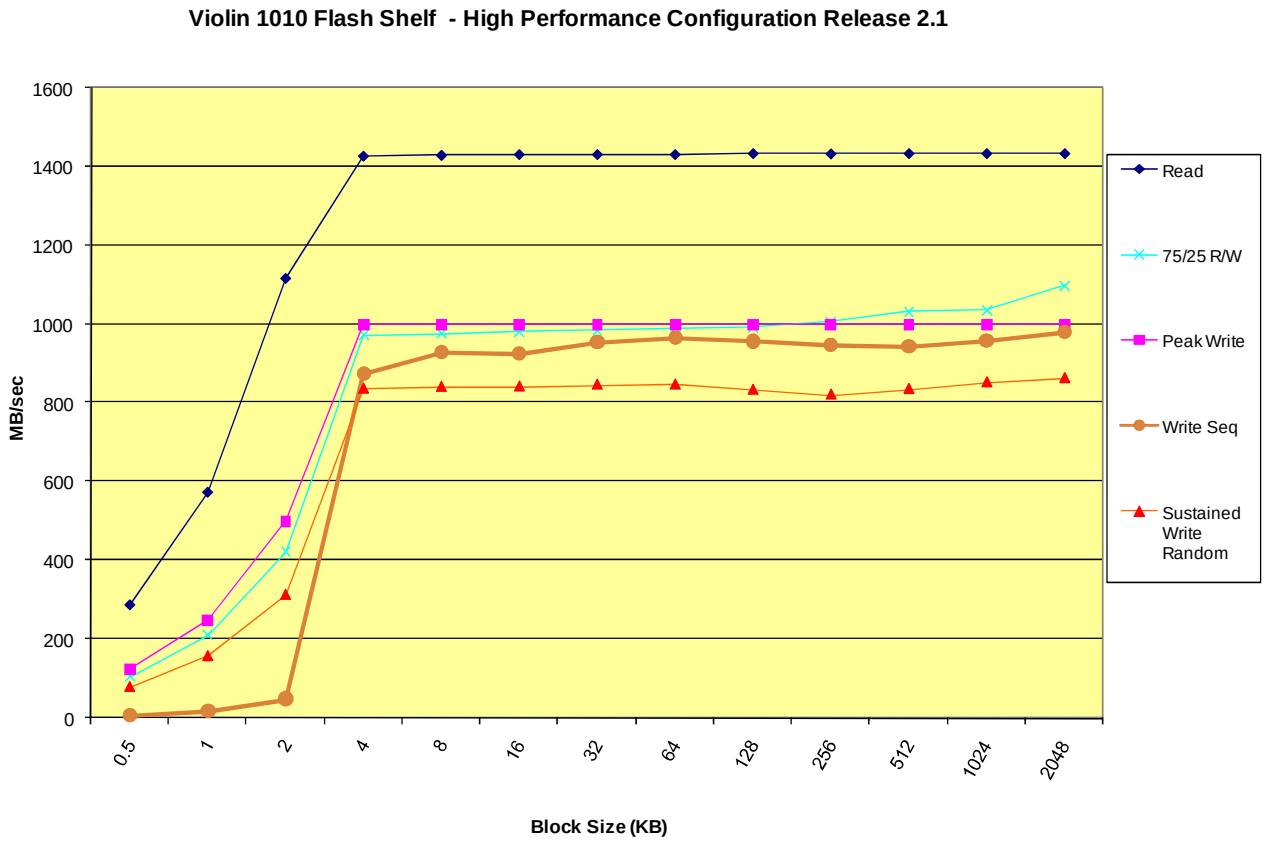

Chart 1: Performance Characteristics of the Violin 1010

The chart shows the Violin 1010 Flash system supports 1 GB/sec of mixed read and write bandwidth. This is sustained bandwidth and can be either random or sequential and equates to approximately 250K IOPS at 4K block size.

Where more bandwidth is required, multiple Violin 1010s can be striped. With the advent of PCIe Gen 2 and the expansion of the PCIe switch chip sets, there are servers that can support over 20 GB/s of bandwidth.

The same approach can be used to scale IOPS. A typical quad-core processor running Oracle can sustain 40K IOPS. Given this, the Violin 1010 is well matched to support 4 or more processors operating at full speed.

Where the Violin 1010 is connected and shared via a network head, the network attachment would typically be dual 10 Gbit/s connections or dual 8 Gbit/s FC connections. This configuration will enable very effective use of the Violin 1010 by many servers and/or clients.

3. SUMMARY

The Violin 1010 Memory Appliance is a high performance solid-state storage unit designed for flexible integration into data centers. Flexible configuration options are available to meet a wide range of capacity and/or performance requirements. Both DRAM and flash memory are supported.

A single Violin 1010 2U appliance with SLC flash memory can support up to 3TB of useable capacity and over 200K sustained IOPS for both random and sequential writes, reads, and read/write workloads.

The most appropriate Violin 1010 Flash configuration is best determined by the system architect after consideration of the performance and capacity requirements of the target application(s).

Violin 1010

Flash VIMM: 32/64 GB SLC

About Violin Memory, Inc.
Violin Memory, Inc. provides the world's fastest and most scalable memory appliances. Violin's products dramatically accelerate large dataset applications demanding very high IOPS or low latency. With support for DRAM and Flash, Violin Memory's Tier-Ø memory appliances provide significant power and space savings with extreme data reliability. Founded in 2005, Violin Memory is headquartered in Iselin, New Jersey. For more information on Violin Memory products, visit www.violin-memory.com.

Product Information and Sales:
Violin Memory
33 Wood Avenue South
3rd Floor
Iselin, NJ 08830
info@violin-memory.com
1 888-9VIOLIN (888-984-6546) x10
www.violin-memory.com/Contact_us