

MACH·2™

Q: How does Exos® 2X14 provide up to 2x the performance of a standard single actuator hard drive?

A: Exos 2X14 can demonstrate up to 2x the performance of a standard single actuator hard drive because it has two independent actuators and data paths, allowing for concurrent I/O streams to and from the host. Each actuator is assigned to a logical unit number (LUN0 and LUN1). For example, one 14TB drive will present itself to the operating system as two 7TB devices, which can then be addressed independently at the operating system as you would with any other hard drive.

Q: What workloads show the best performance benefits over single actuator?

A: Exos 2X14 was designed for hyperscale workloads that focus on low queue-depth random read operations (low queue-depths to keep command latencies low) and large transfer size sequential operations. The highest performance gains over single actuator will be found during high transfer size sequential reads/writes (128KB transfers and larger, queue 2+ for writes), random reads (all transfer sizes), and random writes (128KB transfers and larger).

Q: What performance benchmarks have shown a good benefit over single actuator?

A: Microsoft Jetstress, sequential throughput testing (simulating backup operations or streaming), random read IOPS testing (simulating hyperscale and CDN workloads), and more benchmarks are under investigation.

Q: What types of applications/solutions have shown good benefit over single actuator?

A: Content Delivery Networks (CDN), video streaming, mail servers, backup/shuttle services, Hadoop (benefits demonstrated at a small scale with HDFS and MapReduce), and cloud applications—and more solutions are under investigation.

Q: How are customers testing dual-actuator drives?

A: The best way to test this drive is to use it in production applications so that the true benefits are realized right from the start. Often, synthetic workloads/benchmarks can be confusing and results can be deceiving.

Q: How does Exos 2X14 compare to 10K or 15K enterprise drives in IOPS/TB?

A: Exos 2X14 aims to solve IOPS/TB constraints in systems that are currently using high-capacity 3.5-inch nearline HDDs, whereas 10K/15K mission-critical HDDs have much smaller capacities (2.4TB/900GB top capacities, respectively) and are intended for servers. The 10K/15K solutions are in more direct competition with low-end SSDs, whereas 3.5-inch nearline HDDs are addressing the need for mass amounts of data storage with less-critical IOPS requirements; Exos 2X14 increases the IOPS/TB in these applications. Typically, 3.5-inch HDDs are used in tiered solutions with SSD or NVMe storage solutions to address critical performance needs.

Q: How can I purchase Exos 2X14 drives to improve my IOPS/TB?

A: Contact your Seagate sales representative for more information.

Q: With new technologies like MACH.2™, what is the expected change in power consumption over a single-actuator hard drive?

A: Power consumption will vary by use case, but the highest power mode on Exos 2X14 is sequential reads, followed by random reads. Random reads are the most common access pattern for most of our customers' workloads. Compared to a single-actuator Exos hard drive, when testing random reads at a 4KB transfer size and a queue of 16 pending commands (4KQ16), Exos 2X14 power is approximately 20% higher for up to a 2x gain in performance. Or thought about in a different manner, when considering that you would need to have two single-actuator drives to achieve the random performance of a dual-actuator drive, dual actuator is 40% less power than two single-actuator drives.

Q: What is the typical power capacity of slots in a data center? Is 13.5W somewhere near that number?

A: Customer feedback has revealed a wide range of power capacity in data center slots, but 13.5W is a common max power target that many customers have noted. We look at the aggregate of all the drives to determine the total power you see at a system level. This power limitation is typically more of a JBOD-level limitation where many drives are populated in the same storage system. Many customers are not slot-limited by overall wattage, but rather by 5V power. Exos 2X14 has ~1.4A max current on 5V.

Q: Are there any options to reduce the power consumption of a dual-actuator drive?

A: The on-demand PowerBalance™ feature is available on Exos 2X14, which will detune seek performance and decrease sequential read transfer bandwidth, lowering the overall power read power below 12W. Sequential read power is the highest power mode on dual actuator, but note that many systems may be bandwidth limited by the HBA and system (unable to maintain 500MB/s+ across multiple drives), naturally reducing sequential power consumption.

Q: Why does Exos 2X14 use a SAS interface only with a single port?

A: A SATA interface could be developed on dual actuator, but ATA protocol does not have a concept of LUNs, and thus a different data-placement architecture would need to be developed. It should also be noted that the SATA link speed is 6Gb/s (600MB/s throughput) compared to the SAS link speed of 12Gb/s (1.2GB/s throughput), thus SAS is a better solution to support current and future dual-sequential streaming rates. Initial industry demand has led us to the use of SAS LUNs as an easy plug-and-play solution for currently deployed architectures. Also, the current electrical architecture of the drive only supports single-port as the secondary port is required for linking the SoCs.

Q: Will dual-actuator HDDs ever support dual-port SAS for high-availability and/or failover scenarios?

A: Current electrical architecture cannot support dual-port; this is a design limitation with the current generation of SoCs. Once the market and volume can support a unique SoC, this may be added.

Q: Why is Exos 2X14 only 14TB when Exos X16 is already launched at 16TB?

A: Exos 2X14 is an 8-disk HDD vs the 9-disk Exos X16. The areal density is very similar, but there is currently a 1-disk penalty to incorporate the two actuators.

Q: Can this drive work with software RAID or does it need compatible hardware RAID host bus adapters (HBA)?

A: Yes, this drive can work with software RAID, but considerations need to be made around setting up redundancy across LUNs in the same drive (so that you would not lose data in the rare event of a total drive failure). Currently there are no hardware raid HBAs that fully support dual actuator. Seagate is actively working with industry partners to assure dual-actuator ecosystem support.

Q: What SAS HBAs have been tested with Exos 2X14 drives?

A:

	HBA	9400 Series Host Bus Adapters	<ul style="list-style-type: none"> • Officially supported in P11+ • Driver and controller firmware available on broadcom.com
		9300 Series Host Bus Adapters	<ul style="list-style-type: none"> • Tested and confirmed compatible with multiple LUNs
	RAID	9400 & 9300 MegaRAID	<ul style="list-style-type: none"> • RAID not currently supported
	HBA	1100 Series Host Bus Adapters	<ul style="list-style-type: none"> • Support coming soon (Q1CY20)¹
		2100 Series Smart HBAs	<ul style="list-style-type: none"> • Support coming soon (Q1CY20)¹
	RAID	3100 Series RAID Controllers	<ul style="list-style-type: none"> • RAID-in-HBA-mode support coming soon (Q1CY20) • RAID not currently supported

¹ Estimated schedules subject to change.

Q: What is Seagate doing to assure dual-actuator ecosystem support?

A: We are currently working on a proposal in the T10 committee around command usage and device- versus LUN-level command interaction as a better way for host systems to be aware of potential cross-LUN interactions. We are also actively working with HBA vendors to expand the number of HBA/expander products that are compatible out of the box with Exos 2X14 and future dual-actuator products. Additional socialization of dual-actuator technology is happening at industry events like CES, OCP, Plugfest, and more.

Q: Currently Exos 2X14 is offered in a 4Kn sector size. Is FastFormat™ available to switch to 512e?

A: We plan to offer a 512e configuration in the future, but the initial market demand was for a 4Kn solution. The current version of Exos 2X14 does not utilize FastFormat. Future multi-actuator solutions may implement the FastFormat feature to allow conversion between 4Kn and 512e but in general we have found that customers either use 512e or 4Kn, but not typically both in the same solutions.

Q: Can you highlight the vibration generated and received by using 2 actuators?

A: The system has been tuned so that one actuator can seek with minimal impact on the other. There are no changes to the external drive mechanical specs, such as rotational vibration, shock, acoustic-disturbance, etc.

Q: If one LUN has failed, will the drive still report Test Unit Ready when sent that command?

A: A Test Unit Ready command requires both LUNs to come ready before replying with GOOD status. Note that a LUN failure is very rare, rather head/media errors would be reported at a LUN level as happens on single-actuator drives.

Q: Are there two copies of the defect management list?

A: It is a combined defect list, and it is accessible on both LUNs.

Q: Is a single-volume presented drive in future scope (i.e., the drive itself will load balance/optimize between either the first or second half of the drive)?

A: Maybe; it is possible. There are tail-latency issues with a configuration like this, but it is the simplest way to get plug-and-play performance. There needs to be enough market to justify the firmware development/complexity.

Q: SAS/SCSI command implementation like FORMAT UNIT, which affects both LUNs of the drive, seems like it may cause problems with RAID5, for example RAID5.

A: That is true. We are currently proposing methods in T10 that would address how LUN-versus device-level command interactions are communicated to the system/host. We are considering implementing a command interlock in the future to guard against unintended FORMAT/STOP/etc. commands, but this is dependent on standards adoption/finalization. There are additional RAID complications that could be caused by building redundancy across LUNs without assuring redundancy on additional HDDs.



[seagate.com](https://www.seagate.com)

© 2019 Seagate Technology LLC. All rights reserved. Seagate, Seagate Technology, and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Exos, FastFormat, MACH.2, the MACH.2 logo, and PowerBalance are either trademarks or registered trademarks of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Actual data rates may vary depending on operating environment and other factors, such as chosen interface and drive capacity. Seagate reserves the right to change, without notice, product offerings or specifications. SC702.1-1910US, October 2019