whipTail™
tech | Racerunner™
A Solid-State SAN Appliance

Accelerating Performance with Racerunner Solid-State SANs

White Paper

Contents

Whiptail Technologies
60 Morris & Essex Turnpike
Summit, New Jersey 07901
U.S.A.

Accelerating Performance with Racerunner™ Solid-State SANs

April 5, 2010

Introduction

We live in a brave new digitally connected world, where yesterday's mainframe computers are today's PCs, cell phones, and handheld devices. Nearly every month sees unprecedented breakthroughs in computing hardware and software innovations that were previously thought to be unattainable.

While this technological revolution has provided consumers and businesses alike with devices previously relegated to science fiction, it has also created a performance gap that has seen hard drive-based storage systems lag behind advancements in Random Access Memory (RAM) and processors. As this gap continues to widen, it has allowed solid-state storage to gain more attention as a mainstream storage alternative.

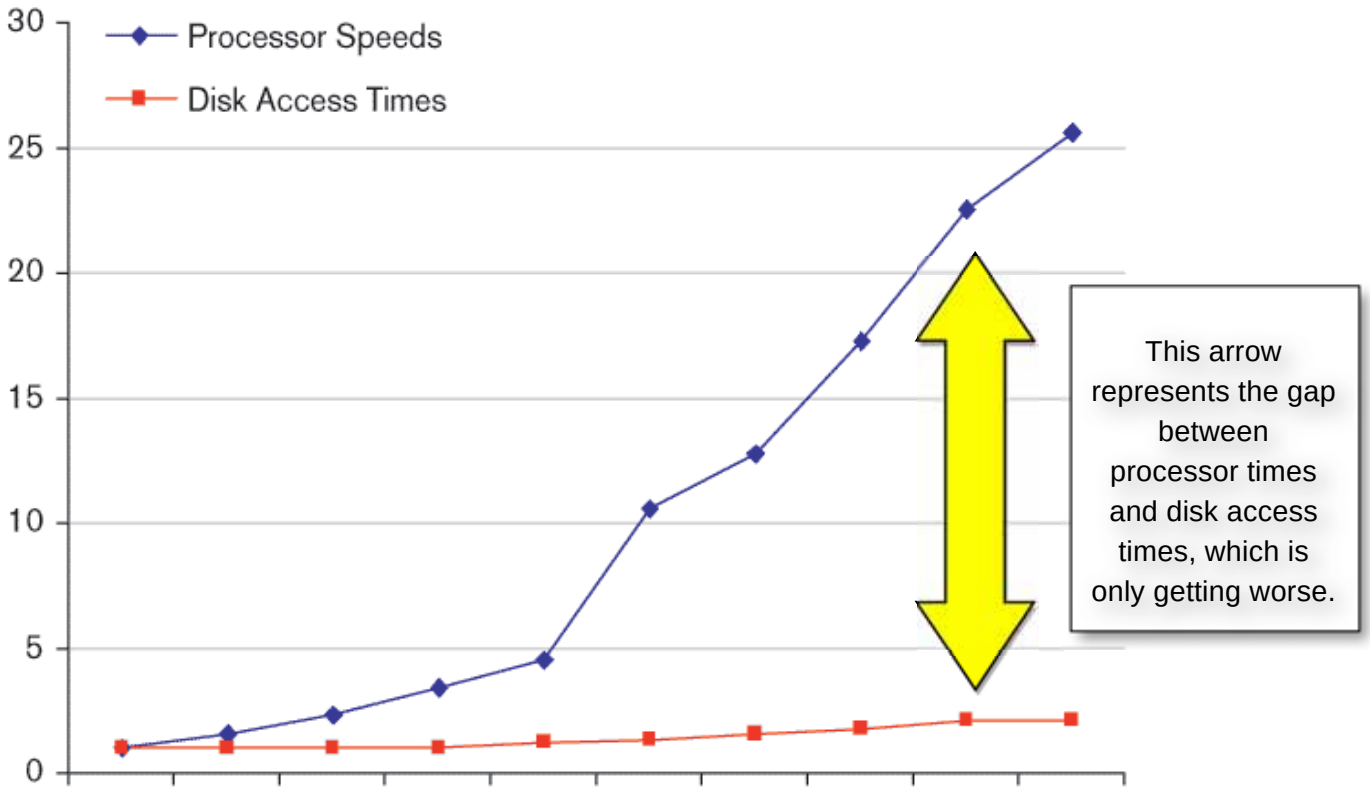

Instead of adopting the mechanical components used by traditional hard disk drives (HDDs), solid-state drives (SSDs) rely on memory chips for data storage. With memory speeds measured in billionths of a second (nanoseconds) compared to disk speeds rated in thousandths of a second (milliseconds), SSDs offer unprecedented access times that narrow the gap between processor and storage speeds.

In addition to super-fast access to data, SSDs provide reduced power usage and higher reliability over conventional HDDs. With these advantages, SSDs deliver on the promise of providing leap-ahead levels of performance for the demanding 24/7/365 applications of today and tomorrow.

This paper describes the benefits of SSDs. It begins with an explanation about why SSDs are needed in today's marketplace, and then goes on to describe their key benefits. This paper concludes with a description of how the Racerunner™ solid-state SAN from WhipTail Technologies lets you get even more out of SSD technology to achieve unparalleled performance and maximize total cost of ownership (TCO) and return on investment (ROI).

Why SSDs? Why Now?

In a 1965 magazine article, Gordon Moore, co-founder of Intel Corporation, observed that the number of transistors per square inch on a microprocessor chip had doubled every year since the integrated circuit was invented. This led Moore to predict that the number of transistors on a chip would double every 18 months — a time interval he later revised to every two years.

While Moore's observation more than 40 years ago remains generally accurate, the curve has not been linear. The formerly marginal rate of advances in central processing unit (CPU) and RAM speeds has increased dramatically over the past decades.

At the same time, one glaring omission to Moore's law is hard disk drive speed, which has not followed the same exponential growth as CPUs and RAM. The gap created by CPU and RAM performance exponentially outpacing hard disk drive I/O and storage performance has become the data center's largest operational constraint. As Intel and AMD continue to break new barriers in CPU technology, servers today are becoming increasingly starved for data. As a result, enterprises are experiencing the effects of disk contention and latency, resulting in large deficiencies in

SSD Trends and Future Direction

Unlike other technologies that promise more than they can deliver and never see the light of day, SSDs are the real deal. In a recent survey, 34% of respondents say they use some form of SSD storage, with 11% having used SSDs for more than two years. While average installations add about 1.5 TB of storage, 80% of SSD users are satisfied or very satisfied with their speedy drives (Source: *SearchStorage.com*).

Moreover, Framingham, Mass.-based IDC predicts worldwide revenue for enterprise SSD storage will increase at a compound annual growth rate of 78% from 2009 to 2012. IDC also predicts that SSD sales will hit $382 million in 2010, up 132% over 2009 (Source: *Storage Technology News*). As SSD prices fall, capacities grow, and storage vendors optimize their arrays and management software to take advantage of SSDs, the nascent SSD market is taking hold and delivering on its promise of becoming a strategic role in enterprise storage.

database and virtualization storage performance.

To compensate, enterprises have been forced to buy (and overbuy) expensive storage in hopes that over-provisioning will yield incremental gains in performance. In the end, the gains achieved are typically measured in single-digit percentages.

From this exercise, it is clear that the magnetic-spinning HDDs that have served the enterprise so valiantly for so long have finally maxed-out on their speed, performance, and ability to keep pace with today's enterprise requirements. For this reason, more and more businesses are turning to SSDs as their storage medium of choice for handling their demanding applications.

SSD's Compelling Advantages

An SSD is a high-performance plug-and-play storage device that uses the same interfaces and form factors as an HDD, but contains no moving parts. Instead of employing mechanical components like traditional HDDs, CDs, and DVDs, most SSDs today are based on NAND Flash chips, a microprocessor-based controller, and software that incorporate these components into a peripheral storage system that acts like an HDD.

The following sections summarize a few of the key advantages that SSDs have over traditional HDDs.

No Moving Parts Mean Faster and Consistent Access Times

HDDs consist of a spinning aluminum disk and an actuator that controls a set of read/write heads. The law of physics associated with these mechanical components limits the rotational speed at which the HDD's actuator can move. Due to issues of stability, power limits, and edge heating, HDDs are limited to about 15,000 rotations per minute, which impose "rotational latencies" of between 2 and 10 milliseconds. This means that an HDD can wait up to 10 milliseconds between when it receives a request and when it starts executing the request.

An SSD, on the other hand, retrieves information from a NAND cell in about the same time it takes a server to retrieve data from RAM, typically 0.1millisecond. The only performance restrictions placed on SSDs are limitations imposed by quantum-level mechanics at the chip level.

Because physical location of data is irrelevant for SSDs, they do not have to rely on a read/write interface head to synchronize with a rotating disk, thereby making read performance more consistent than HDDs. And with no platters to spin or heads to move, SSDs offer faster boot and application launch times.

Although SSD read and write operations outperform HDDs, their random read performance is faster than their random write performance. This performance gap is due to a collection of bytes called an "erase block" inside the drive. When you write to an SSD, the drive merges

your changes with existing data to update a complete erase block. These erase blocks can cause random write operations to be nearly three times faster than SSD read operations.

Regardless of this imbalance between SSD read and write performance, an SSD-based RAID set easily outperforms an equivalent number of traditional HDDs in the same configuration. In tests comparing the read performance for random data sets, at least four times the number of HDDs may be required to match the performance of an equivalent SSD configuration. This vast reduction in the number of drives lets enterprises reduce the number of high-performance drives required to meet given performance levels, resulting in significant cost savings, about 90% when compared to traditional storage costs. These savings come from not only the decreased amount of drives needed but also the dramatic reduction in power, cooling, retail rack space, and overall maintenance required.

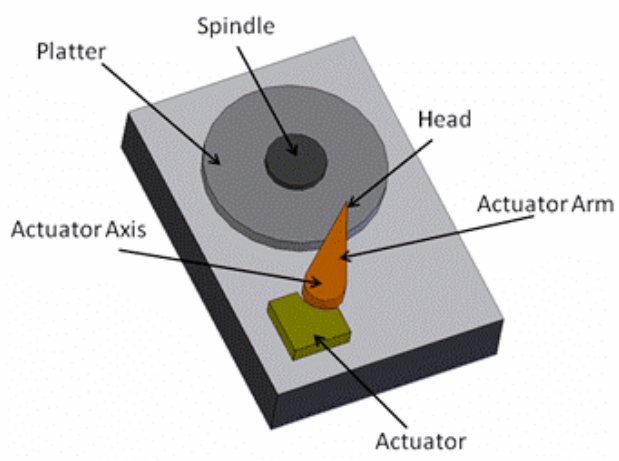

Superior Longevity and Reliability

The mechanical components in HDDs are susceptible to wear and tear. Even under the best of circumstances, within an ideal operating environment of an air conditioned office or data center, HDDs deteriorate over time simply from normal wear and tear on the moving recording mechanisms each time they are turned on and read from or write to the disk.

The frequency of drive failures, crashed drives, and lost data are well documented. Horror stories abound about how valuable data was lost and the anguish of trying to retrieve a portion of a crashed hard drive. Entire industries and submarkets have arisen to help safeguard business entities from the all too common problems with HDDs offering back up, archival, off-site storage, and data-recovery services.

Risking your data to a fragile recording mechanism that can fail when you need it most does not make sense when dependable SSDs are now available.

SSDs also provide the ruggedness required in demanding environments. Unlike HDDs, whose fragile and sensitive components can be damaged by small jarring movements or harsh environments that can render the drive completely unreadable, SSDs can withstand

extreme shock and vibration with data integrity, and deliver reliable performance in extreme conditions.

In the end, no matter how much protection is provided for conventional HDDs, no matter how buffered the drive enclosure is from shock and vibration, none of that addresses the primary flaw in the HDD's design: that the drive starts wearing out the day it is first turned on and with each hour of data it reads and writes, it gets closer to failure due to natural wear and tear.

Reductions in Power Consumption and Heat Dissipation

With no platters to spin or heads to move, SSDs consume about a third of the power of HDDs. Along with the lower power consumption, SSDs have considerably less heat dissipation than HDDs due to the absence of heat generated from the rotating/movable media. With the absence of cooling fans, SSDs run completely silent compared to the spin of a magnetic HDD.

Applications That Benefit from SSD Deployment

SSDs are an excellent solution for I/O bottlenecks, particularly those caused by the high access times of traditional HDDs.

SSDs are most effective for server applications and server systems, where I/O response time is crucial. Data stored on SSDs should include anything that creates bottlenecks, such as databases, swap files, library and index files, and authorization and login information. Examples of such applications ideally suited for SSDs include:

- Virtual desktops (VDIs)
- Server virtualization
- Relational databases
- Enterprise messaging
- Metadata indexing and logs
- Data warehousing applications
- Online transaction processing (OLTP)
- High-speed data acquisition
- Multi-tasking systems that use high-speed swap files

By delivering unprecedented access times, SSDs empower enterprises to reach performance levels they could never approach with traditional HDDs. As SSD capacity increases and prices continue to fall, enterprises running applications that need certain

levels of performance can achieve greater productivity with SSDs at a more cost-effective price point, with a superior ROI.

The Racerunner™ Advantage

The previous sections in this paper described how SSDs deliver the highest levels of performance, data integrity, reliability, and durability. The Racerunner™ solid-state SAN from WhipTail Technologies is a patent-pending 2U SSD storage-area network (SAN) appliance that leverages these SSD advantages and more to remove the system performance bottleneck in enterprise storage applications.

The Racerunner contains an array of solid-state drives overlaid with expert software to create a powerful primary storage solution that accelerates mission-critical data and dramatically expands virtual capacity. Racerunner delivers an incredible 150,000 IOPS of read performance and 65,000 IOPS write performance, with 0.1 ms latency. Racerunner is available in 1.5, 3.0, 6.0, and 7.5 TB models. All models include a redundant power supply and can be set up in RAID 0, 5, 6, and 10 configurations.

Racerunner is compatible with all the major storage protocols, including:

- Fibre Channel
- Internet Small Computer System Interface (iSCSI)
- Infiniband
- Network File System (NFS)
- Common Internet File System (CIFS)

As mentioned earlier, SSDs are inherently more reliable than HDDs because their data path does not require mechanical arms, spinning platters, or other moving parts. To increase their operational longevity, Racerunner utilizes proprietary wear-leveling technologies embedded on the drives that evenly distribute data across all of the available storage space during write operations. This eliminates the constant re-use of one small portion of the storage medium and extends the lifetime storage of your valuable data to beyond seven years, dramatically optimizing the potency of the solid-state SAN's underlying Flash architecture.

Deduplication and Compression

In addition, Racerunner is the only appliance on the market to offer both deduplication and compression primary storage “in-line” and “on-the-fly,” without having to store duplicate data.

- Data deduplication is a way of reducing redundant data. While other companies offer a similar technology to reduce backup or recovery storage, WhipTail Technologies is the only vendor offering inline data deduplication where data redundancy is eliminated on the ingress at the block level. This approach is truly the “holy grail” of deduping because it dramatically decreases the amount of storage an enterprise requires. This benefit is a huge advantage in virtualized environments and critical for establishing Tier 0 performance as a viable solution for the enterprise.
- In-line data deduplication is a standard feature of Racerunner that can be customized by each enterprise to suit their requirements by turning in-line deduping on a per-LUN basis on or off at will. Whether your environment requires optimized virtualization, accelerated data, or toggling between the two, Racerunner provides the optimal business flexibility and performance to meet your changing needs.

The following real-world reduction ratios identify how effective these deduplication and compression features can be:

- Virtualization environments achieve reduction ratios of up to 300:1.
- Database and network-attached storage (NAS) environments achieve typical reduction ratios of 4:1.

Moreover, Racerunner removes up to 90% of the Tier One storage currently provisioned for VMware® ESX, Xen, and Microsoft® virtual machines. For more information about Racerunner’s real-world virtualization and database successes, see “Applications That Benefit from Racerunner” on page 10.

Note: All deduplication features are not cut from the same cloth. For example, some Tier One vendor storage arrays that claim to perform in-line deduping actually increase, not decrease, the amount of storage required and introduce severe performance degradations. The reason is two-fold:

1. The deduplication algorithms used commit all write operations to physical media during production hours. This requires sufficient space to hold all customer data on expensive storage prior to reduction.
2. At scheduled times, the filer fingerprints all new data, compares it to the disk fingerprint table, and performs deduplication. This process is performed as a batch job and is computationally intense, which reduces performance during the entire deduping process.

Thinking Green

To reflect WhipTail Technologies' commitment to protecting the planet, Racerunner offers environmentally friendly innovations that deliver lower temperatures and quieter operation. These "green" features make the array suited for a variety of eco-friendly storage usage models, especially those with tighter power-consumption constraints. And with its outstanding performance, data centers can replace their fully populated storage shelves with eco-friendly 2U Racerunner arrays at a 1:30 ratio, reducing power consumption and space by as much as 90%.

Best of all, Racerunner is truly plug and play. Plug it into a power source of less than 200 Watts, assign a network protocol and logical unit numbers (LUNs), and you're ready for business.

With all these features, Racerunner makes Tier 0 performance a realistic option as a primary SAN for today's enterprises.

Applications That Benefit from Racerunner

While nearly all applications can benefit significantly from Racerunner's outstanding performance, nowhere is this impact felt more, and the ROI and TCO greater, than with virtualization and database applications. Running VDI and database applications on high-performance, energy-efficient Racerunner arrays slashes disk contention and access times, allowing virtualization and database servers to process more data in dramatically less time. Internal and independent tests show more than a 95 percent reduction in database report generation and load time compared with traditional magnetic disks.

The following sections provide greater detail about the benefits that can be achieved by running virtualization and database applications on Racerunner arrays.

How Racerunner Scales Virtual Desktops

Virtual Desktops (VDIs) have more in common than just virtualization. They all require an overwhelming storage infrastructure to support their demanding computing architectures. While the paradigm shift of hosted desktops, cloud-desktops, and other VDI environments is finding acceptance in user communities, the absence of high-performance storage solutions that can unleash VDI to provide acceptable user experiences is preventing data centers from experiencing the virtualization benefits that the larger user community is enjoying.

Consider the following:

- A virtual desktop requires from 20 to 40 IOPS per image to run at an acceptable level.
- Fibre Channel HDD arrays from common Tier One storage companies such as EMC and NetApp deliver approximately 200 IOPS per drive.

This means a 5,000 person user environment requires roughly 250,000 IOPS to perform at minimal user expectations and standards. At 200 IOPS per HDD, 625 drives are needed to meet this performance requirement. This corresponds to a performance cost ratio of $8 per I/O operation per second. Using Racerunner, this cost drops to 37 cents per I/O operation per second!

As if this cost differential was not staggering enough, consider that the total cost of a 625-drive Fibre Channel storage array, including drives, enclosures, and controllers, etc. would exceed seven figures upfront and consume more than 60 units of data center rack space. As a result, the combination of excessively high initial CAPEX expenditures and pricey ongoing costs for power, cooling, and rack space cost makes it difficult to justify ROI, especially when organizations are trying to improve their bottom line. For this reason, many VDI projects never see the light of day, as they simply cannot be business-justified.

When solid-state drives are offered by the primary Tier One HDD manufacturers, they land in a price-point upwards of $35,000 per drive that delivers only 6,000 IOPS each. Their inability to master the nuances of SSDs by adding shelf after shelf of expensive DDSs does not add to performance, which normally maxes out at around 60,000 IOPS. With an investment cost of $350,000 and without seeing a dramatic performance increase, the cost is about $6 per I/O operation per second making this another no-go for most enterprises today.

There is a bright side to all of this and its name is Racerunner. Racerunner delivers the best cost per I/O operation on the market, with an industry-changing, VDI-enabling cost of just $0.37 per I/O. Available as a 2-unit storage appliance, in your choice of capacities (1.5 TB, 3.0 TB, and 6.0 TB,), Racerunner boasts an all-inclusive MSRP starting at $49,000. Move over, Tier One storage vendors, and welcome to the new world order!

The Challenge with Virtualization

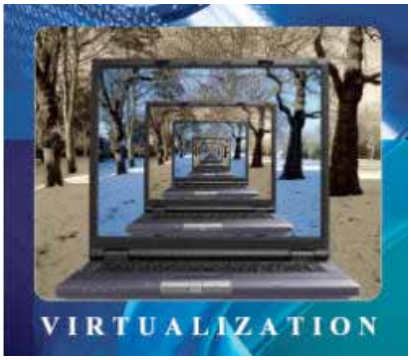

Virtualization environments are typically over-subscribed when it comes to disk performance, with everyone competing for the same pool of resources. This is typically not a problem in an environment where low utilization servers have been virtualized. However, when mid-to-large organizations try to virtualize tier one applications (for example, Microsoft Exchange, etc.) major disk performance issues often emerge. For example, oversubscribed disk systems are notoriously

hard to identify. In addition, CPU and memory usage is much more “in the face” of administrators. Disk systems require intimate knowledge of the underlying disk infrastructure and the workloads running on them. Virtualization can make this even more challenging, as most monitoring and management tools only show you the amount of data being transferred back and forth.

With many systems competing for the same disk resource, SEQUENTIAL workload can rapidly turn into much more demanding RANDOM workloads. As it turns out, RANDOM workloads are the worst possible workload for mechanical disks, because with virtualization, every virtual machine gets a timeslice of every available resource (CPU, memory, disk, I/O, and so on). As a result, a nice sequential file transfer, such as a large ISO image, gets interrupted every so many milliseconds by another fileserver or an Exchange server, causing the read/write head to move to another track and incurring a seek penalty.

Since the Racerunner solid-state SAN is completely solid-state in nature, average data access time is sub-millisecond (0.1ms.) As a result, SEQUENTIAL and RANDOM workloads are no different to Racerunner.

While virtualization overhead can be seen as low (6-15%), it is still significant and impacts performance for I/O-intensive applications. The Racerunner’s latency advantage, even with 15% overhead, is 30-to-60 times faster than running without a hypervisor. With Racerunner, enterprises can finally recognize the value of virtualizing this tier of application by making the “performance pool” dramatically larger, thereby eliminating the contention.

How Racerunner Excels with Databases

With the dramatic pace of development from chipmakers Intel and AMD, today’s enterprise servers are plain and simply being starved for data. While processing speeds have increased exponentially, the speed at which a disk spins has not changed in over 7 years – a lifetime in the world of computers. This performance gap is widening with each release of high-performance processors.

The underlying problem is that database workloads are extremely demanding for traditional spinning disk systems. The majority of the traffic, however, is based on extremely small request sizes and is often random in nature. As a result, an HDD’s mechanical latency becomes a limiting factor in database performance. Barring a cache “hit” (data already being in a RAM buffer), each request requires the drive’s read/write head to move, resulting in a 6 to 9 ms delay while the head moves and the data is retrieved.

When you add the speed limitations of traditional drives being 15,000 RPM, you realize that HDDs can only sustain operations of 200 IOPS at best. It becomes clear that customers need to over-provision in HDDs to gain a minor boost in performance. This vicious, and costly, circle is one of the biggest challenges facing enterprises today, and is directly responsible for the huge percentage of an enterprise’s IT budget – in many cases over 60% – being allocated solely for storage.

The solid-state Racerunner is specifically tuned for low-latency, high-I/O environments. Its horsepower more than accommodates the demanding requests of today's high-performance processors (including Nehalem and Shanghai CPUs). With its proven ability to handle over 150,000 I/O operations per second at small request sizes (4K) with submillisecond latency, the Racerunner can outperform other arrays without breaking into a sweat. And with its minimal data-center footprint of 2U and limited power consumption (<200 watts), the Racerunner is the perfect strategic solution for empowering IT departments to get ahead of the curve when deploying new applications and projects.

Reducing Storage for VMware ESX and Xen Servers

Racerunner's inline data deduplication/compression features free you to go way beyond thin provisioning. By eliminating the need to store repetitive data, there is no need to continue wasting precious Tier One storage by supporting hundreds of copies of your virtual machine's operating system (usually Microsoft Windows 2003, 2007, etc.)

With Racerunner, you can dramatically optimize your storage by using only enough space to store 10-15% of what you have provisioned. Since Racerunner's inline deduping engine is truly "on-the-fly" and not run in backup or as a batch-job during off-peak hours, you never have to store duplicate data again. Now companies can extend the value of their current investment in HDDs simply by re-allocating the multiple terabytes of freed-up spinning disks to alternative databases or environments.

In parallel, our customers achieve additional consolidation ratios thanks to Racerunner's compression capabilities. Although databases do not usually suffer from heavy doses of repetitive data, they are alphanumeric and compress extremely well, thus compression ratios of 4:1 are common.

About WhipTail Technologies

WhipTail Technologies is a SAN solid-state storage manufacturer located in Summit, New Jersey. In recognition of our high-performance SAN solutions, the company is named after the WhipTail Racerunner lizard, an extremely fast species indigenous to the southwestern United States. Racerunner SSD SAN array was released to general availability in late 2008.

At WhipTail Technologies, we know that when a vendor touts breakthrough technologies, it has to mean something. The industry is competitive, and it's not good enough to offer vague platitudes and the same promises everyone else is making. For this reason, WhipTail Technologies backs up its words with a 30-day corporate evaluation program that allows companies to see how Racerunner can make a dramatic difference in their workplace. To share our excitement and see what Racerunner can do for you, please contact us by phone at 1(888) 550-8136 or visit our Web site at www.whiptailtech.com.