


## Scalable Informatics JackRabbit

A High Performance, High Density, Scalable Storage System



## Introduction:

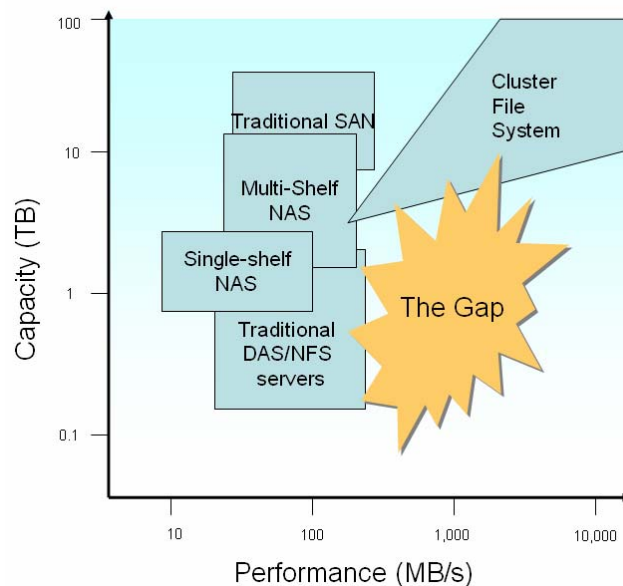
Scalable Informatics JackRabbit systems are multiprocessor storage servers providing up to 36 TB of fast, but low cost storage in support of solving today's most complex computationally and data intensive problems. Scalable Informatics JackRabbit has demonstrated over 1 GB per second performance to internal RAID6 disks. Scalable Informatics JackRabbits are capable of providing multiple GB/s over high performance networks. Scalable Informatics JackRabbit supports a full range of storage subsystems including RAID 0, 1, 5, 6, 10, 50 and 60 with hot-swappable SATA 300/NCQ disks

Scalable Informatics JackRabbit supports NFS or CIFS/SMB file systems, FTP file servers, and iSCSI target and initiator and is easily integrated into storage management systems. Scalable Informatics JackRabbit currently runs Linux and will support other platforms in the future.

## The Challenge:

Modern high performance computing systems, such as Linux clusters for computation, simulation, near-real time data analysis, database operations, rendering or visualization all generate, process and need to store huge amounts of data. Users of these high performance systems require both storage capacity and performance.

A range of different storage system designs exist with varying levels of complexity, cost structures, and performance limitations. Unfortunately most of today's storage systems have failed to address the critical areas of the capacity-performance space, where a cluster file system and SAN would be overkill and cost prohibitive, while a traditional NAS server would be an under-performer at best, and provide insufficient capacity.

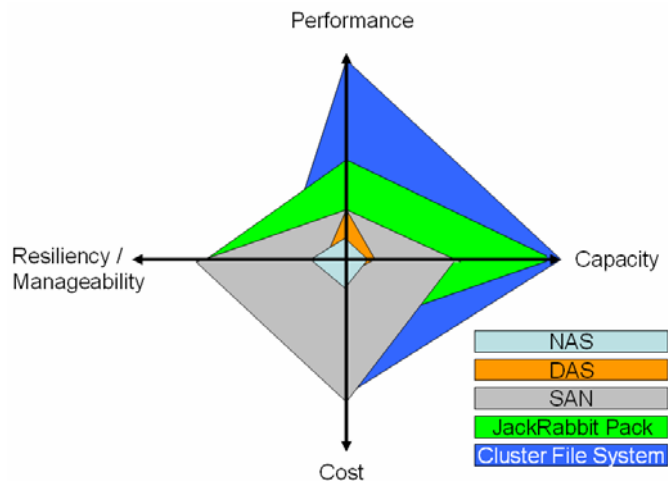


This is where Scalable Informatics JackRabbit excels. Scalable Informatics JackRabbit systems address the critical areas of the performance and capacity space, enabling you to provide very high performance and capacity for your users, without adding complexity in design or management.

## ***Built for performance and reliability***

The Scalable Informatics JackRabbit comes in both a 3U and 5U form factor and both are designed with reliability and performance in mind.

From a reliability perspective, the system has multiple hot-swappable power supplies in an N+1 configuration that provide continuous power to the units AND multiple hot-swappable cooling fan units that provide two distinct airflow paths for disk and system cooling in the 5U unit AND multiple hot-swappable fans provide ample cooling for the 3U unit.



Each 3U unit may be configured with one or more RAID controllers, while each 5U unit may be configured with two or more RAID controllers. With one RAID controller, Scalable Informatics JackRabbit unit provides in excess of 500 MB/s sustained throughput to a RAID6 device as measured by IOzone, bonnie++, and spew. With two RAID controllers, Scalable Informatics JackRabbit provides in excess of 1 GB/s sustained throughput to a RAID6 device.

The multiple RAID cards may be configured as a RAID1 mirror, a RAID0 stripe, or combined with DNBD/DRBD and heartbeat, it may be used as a high availability system.

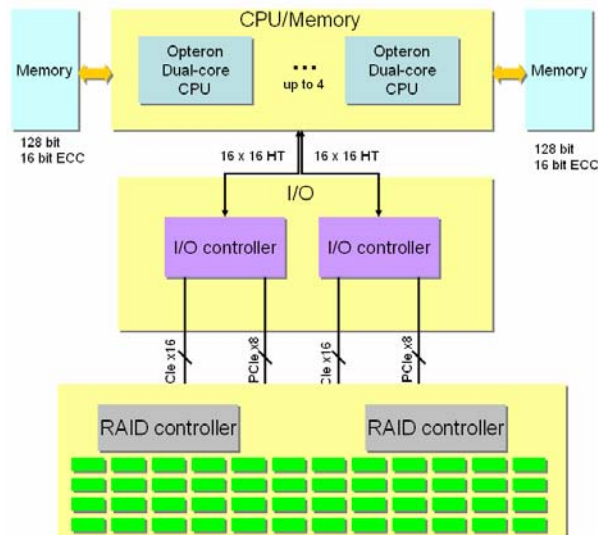
Each Scalable Informatics JackRabbit unit may be configured with 2 or 4 x dual core Opteron processors, providing from 4 to 8 processor cores per unit. Up to 32 GB DDR2/667 RAM is provided in the 2 x dual core system, and up to 64 GB DDR2 ram is provided in the 4 x dual core system. Multiple network connection options include multiple quad gigabit Ethernet, multiple Infiniband, and 10Gb Ethernet adapters. Some adapters may impose specific configuration constraints. Other networking technologies would be considered once they are supported under Linux and qualified to work with the unit.



IPMI 2.0 will be provided for system monitoring and control. Multiple USB2 ports, a serial port running a serial console for remote control, a VGA and PS/2 keyboard and mouse ports are provided.

### **Architected to scale: I/O bandwidth**

Scalable Informatics JackRabbit is designed to scale, with significant available I/O bandwidth. Each JackRabbit unit includes multi PCIe lanes providing in excess of 44 lanes (22 GB/s) of internal IO bandwidth. These PCIe lanes are available for connecting RAID and network interfaces. The 5U Scalable Informatics JackRabbit product has two x16 and two x8 PCIe interfaces. An additional 2 PCIx busses are available, providing 2 x 850 MB/s of bandwidth



Multiple RAID controllers may be used for redundancy or performance. A minimum of two RAID controllers are required for the 5U system. Each RAID controller is connected to the system via a multiple lane PCIe bus. Network I/O may be attached to a PCIx or PCIe bus if required by network bandwidth considerations, such as DDR Infiniband or 10 Gb Ethernet.

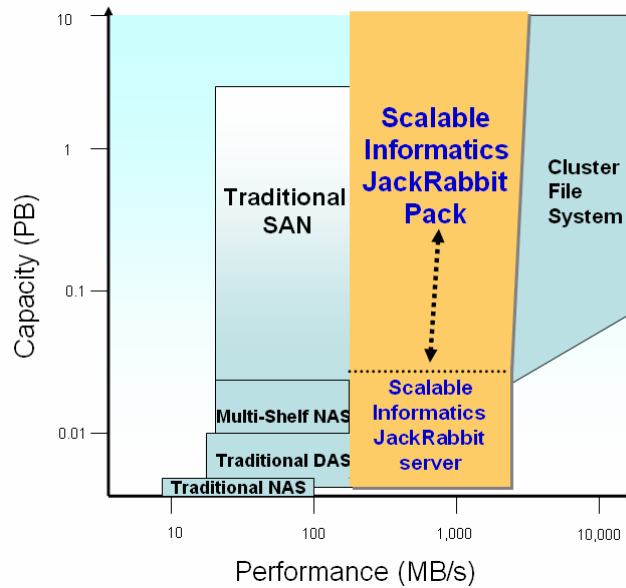
Each unit has a quad Gigabit Ethernet adapter, which may be channel bonded. This provides a default of 400MB/s in network bandwidth. An additional 2 quad Gigabit Ethernet adapters options are available, for a maximum of 1.2 GB/s bandwidth to a Gigabit-class network per chassis. Additional options, namely Infiniband or 10Gb Ethernet may also be selected. If these options are chosen, the adapters would occupy a PCIe slot, unless a PCIx slot version was available.

With a dual port, DDR Infiniband connection, Scalable Informatics JackRabbit would have approximately 4 GB/s of network bandwidth. With a single port 10 Gb Ethernet card, Scalable Informatics JackRabbit would have approximately 1 GB/s of network bandwidth.

### **Scalability:**


Multiple Scalable Informatics JackRabbit servers may be linked together using DNBD/DRBD to create high availability units if used with heartbeat, STONITH, and other technologies. Scalable Informatics JackRabbit servers can export iSCSI volumes, and can be aggregated behind an iSCSI initiator.

Using an Infiniband switch within an data cluster, multiple Scalable Informatics JackRabbit servers may be aggregated at very high bandwidth per unit by leveraging the iSCSI capability of these units. Extremely large file systems could be built in this manner; with 8 Scalable Informatics JackRabbit servers per 41U rack, each rack would be able to provide 288 TB of raw storage, with an Infiniband or gigabit switch to link these units.



With 4 racks (32 Scalable Informatics JackRabbit units), you would be able to provide 1.1 PB raw of data storage. Setting up 2 RAID6 units per chassis, users would be able to provide 248 TB of RAID6 per rack, or 0.992 PB over 4 racks. Using a single iSCSI initiator, users could set up a file system across these units, provided the file system supported these capacities. Both XFS and JFS support this capacity and beyond, under Linux.

To build such a system, a configuration approximately as follows would be used:




This relatively compact system could provide nearly 1 PB of storage in 4 racks, interconnected via a high performance Infiniband fabric. With an appropriate filesystem, and additional volume management technology, this system would be capable of providing performance of multiple GB/s to the end user.

A no-single-point-of-failure design would take into account potential failure modes, and provide protection in terms of redundancy. Multiple technologies exist to provide this capability within Linux. DRBD and heartbeat enable building HA systems in an active-passive mode. Leveraging the Oracle Cluster File System (OCFS2), we can provide a clustered file system across units.

The following block design provides approximately 1 PB of storage, in a high availability/no-single-point-of-failure system, in 9 racks, where each rack consumes less than 15kW.

JackRabbit pairs run DRBD and HA, can perform active-passive or active-active.


Pairs interconnected via redundant Infiniband switches, and heartbeat and failover.



Additional IO bandwidth may be added by “striping” across more of these units. It is also important to note that 1 TB disk units will be available shortly, which will enable the Scalable Informatics JackRabbit to provide 48 TB raw, approximately 44 TB in a RAID6 configuration, and about 42 TB after formatting. 8 units would provide 336 TB of formatted capacity, and 32 units would provide 1.3 PB of formatted storage.

## ***Compatibility:***

Scalable Informatics JackRabbit servers are shipped with OpenFiler 2.x so as to be able to operate as an appliance upon installation. The Scalable Informatics JackRabbit servers support Linux, including SuSE, RedHat, Ubuntu, rPath/OpenFiler, Centos, Debian, and others. Other operating systems may be supported if sufficient customer or partner demand arises, including Solaris 10, \*BSD, and Windows x64 Server OS. You may run any OS you wish upon it. Scalable Informatics will support Linux and Linux variants which provide up to date kernel drivers, and other operating systems as demand warrants.



## ***Filesystems:***

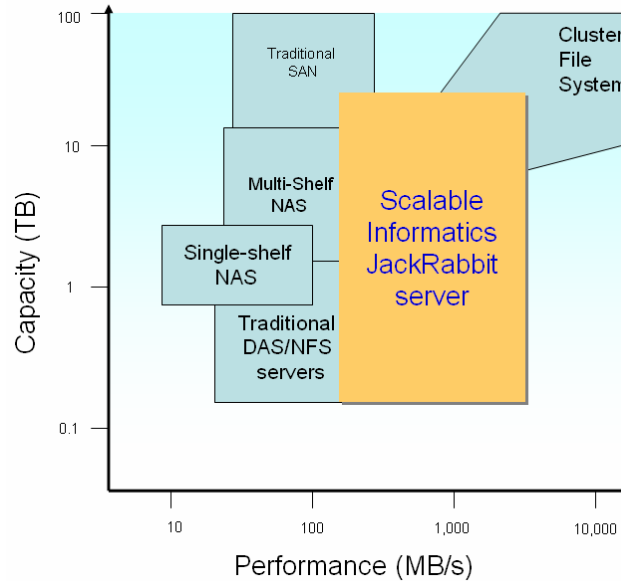
Under Linux, the Scalable Informatics JackRabbit unit will support XFS, JFS, OCFS2, ext3, NFS, CIFS/SMB, and Reiserfs. Some distributions may not support some of these file systems without patching. All Linux distributions will support iSCSI. Under Solaris 10, zfs may be supported. Under Windows x64, NTFS and CIFS/SMB will be supported.

## ***Support:***

Scalable Informatics provides a 3 year parts replacement warranty on JackRabbit. Additional levels including next day parts replacement, on-site support are available as extra cost options.

## ***Summary:***

Scalable Informatics JackRabbit series systems address critical performance and capacity requirements for end users whose applications may not require the performance of a cluster file system, or the complexity of a SAN. These systems provide tremendous performance, capacity, configurability, and unbeatable price performance.



Scalable Informatics JackRabbit systems are designed to fit easily into your environment, interoperating with your systems. They are designed to be easy to manage, and simple and cost effective to support. They address the capacity-performance gap between NAS/DAS/SAN and Cluster File Systems.

Using Lustre or OCFS2 you may build or incorporate Scalable Informatics JackRabbit into Cluster File Systems. The Scalable Informatics JackRabbit could run PanFS from Panasas, and provide a connection into highly scalable file systems, providing disk-to-disk backup capability, as well as a RAID6 based home directory for such clusters.

The Scalable Informatics JackRabbit may be used for scalable database clusters, running Oracle RAC, MySQL, PostgreSQL, and other database systems. Multiple Scalable Informatics JackRabbit units may be configured as shared-nothing with replication, or active-passive using DRBD, or even active-active using OCFS2 or similar products.

The Scalable Informatics JackRabbit may be used for disk-to-disk backup storage, or as a local nearline cache and attach point for offline storage (tape, WORM, and related technologies).

The Scalable Informatics JackRabbit server is a highly flexible server product, able to take on some of the most challenging problems, and address them in a high performance and cost effective manner.

For more information, please contact us at

[sales@scalableinformatics.com](mailto:sales@scalableinformatics.com)

<http://www.scalableinformatics.com>

voice: +1 734 786 8423

fax : +1 866 888 3112