



# PROVEN SCALABILITY FOR SAS 9.2

SCALE UP. CONSOLIDATE. SAVE.

Today, companies are increasingly utilizing analytics to discover new revenue and cost-saving opportunities. Many business professionals turn to SAS, a leader in business analytics software and service, to help them improve performance and make better decisions faster. Analytics are also being employed in risk management, fraud detection, life sciences, sports, and many more emerging markets. However, to maximize the value to the business, analytics solutions need to be deployed quickly and cost-effectively, while also providing the ability to readily scale without degrading performance. Of course, in today's demanding environments, where budgets are still shrinking and mandates to reduce carbon footprints are growing, the solution must deliver excellent hardware utilization, power efficiency, and ROI.

To help solve these challenges, Red Hat, SAS, NEC, and Intel collaborated to prove the linear scalability of SAS 9.2 running Red Hat® Enterprise Linux® on NEC's newest Intel® Xeon® processor 7500 series-based platform. The result is a pre-tested, scalable server configuration that can help you deploy faster, reduce risk, lower cost, and plan for future upgrades.

## SAS 9.2

SAS 9.2 provides the core components of the SAS Business Analytics Framework and significant performance improvements over SAS 9.1 on Linux. SAS 9.2 helps users gain insights that are often hidden in data, so they can reach evidence-based decisions with confidence. SAS 9.2 supports the entire analysis process – from data access to the point of decision – however varied or complex. A wide range of data integration techniques empowers users to collect, classify, process, analyze, and interpret data to reveal new insights. SAS 9.2 advances the capabilities of SAS analytical products, including forecasting, data mining, optimization, and model management. SAS

Analytics provide rapid answers to key business questions, allowing decision makers to react more quickly to fast changing conditions.

SAS 9.2 is available as 64-bit enabled applications supporting 64-bit extended architectures. This enables you to scale up or consolidate multiple SAS instances within one affordable, powerful, commodity system.

*“With the arrival of a new generation of industry standard SMP servers with 48, 64, and even greater core counts, we've had many customers inquire about how SAS Analytics will scale up with Linux on these new servers. These tests on the NEC platform have proven to us that Red Hat Enterprise Linux can scale vertically to exploit the potential of these servers while also delivering the intense I/O throughput that is characterized by SAS Analytics.”*

Craig Rubendall

SAS Director of Research and Development

## PROVEN SCALABILITY

SAS created multi-user benchmarking scenarios to simulate the workload of a typical SAS Foundation environment. The goal of these scenarios is to evaluate the multi-user performance of SAS 9.2 on various platforms. Two tests are used:

### CPU-focused

- Heavily computational with very light I/O with a goal to fully utilize CPU power
- Row counts up to 187,000, under 10 GB total I/O data

### Mixed

- A mix of CPU and I/O-intensive jobs with a mix of short and longer running jobs
- Row counts up to 90 million, variable counts up to 297, and over 800 GB total I/O data

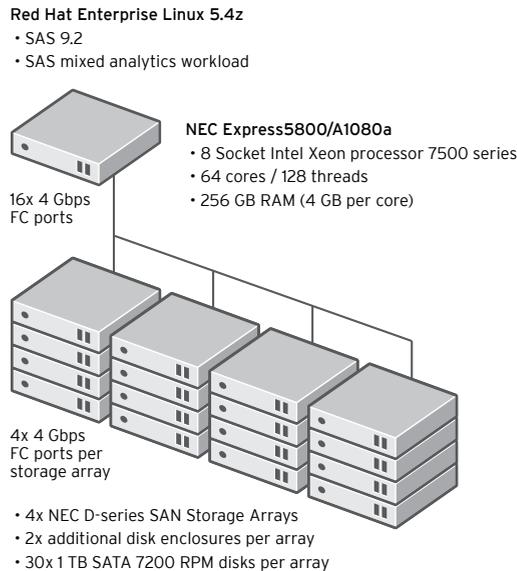
IN COLLABORATION WITH:



NEC



**LINEAR SCALABILITY TO 64 CORES.  
CONSOLIDATE WORKLOADS. BETTER UTILIZATION.  
GREENER SAS DATACENTER. SAVE.**



Tests were conducted with 32 cores and 500 GB of input data to provide a baseline, and then with 64 cores and 1 TB of data to determine scalability. The results were impressive, demonstrating near linear scalability when the load and resources were doubled. The following features contributed to the performance and scalability of the configuration:

- Intel Advanced Programmable Interrupt Controller (APIC) for optimized Non-Uniform Memory Access (NUMA)
- Support for Red Hat Enterprise Linux on the new multisocket NEC Express5800/A1080a server
- NEC NUMA aware BIOS – significant performance gains
- Red Hat Enterprise Linux processor scheduler – automatically optimizes SAS application processes

- Intel Turbo Boost Technology and Intel Hyper-Threading Technology – improved scalability

The Red Hat Enterprise Linux tunable I/O stack was essential in increasing SAS 9.2 performance – only standard best practices for system tuning were applied.

- Read-ahead on Logical Volume Manager (LVM) devices was adjusted to 8,192 bytes
- Standard blockdev tool was utilized to accommodate large sequential access to LVM devices and the file system
- Performance was increased by 30 percent by using the XFS file system instead of ext3 (XFS is better suited for the large sequential I/O that is typical in analytics workloads)

**SCALE-UP, CONSOLIDATE, AND SAVE WITH RED HAT ENTERPRISE LINUX ADVANCED PLATFORM**

Red Hat Enterprise Linux provides an easy, no-risk path to open source cost savings and performance. Basing your solution on open source technologies wherever practical helps you make the most of your IT budget and take advantage of the latest technology innovations. In addition, Red Hat's robust ecosystem includes thousands of software and hardware vendors actively engaged in supporting and certifying Red Hat Enterprise Linux, often before other operating system platforms.

**Vertical performance.** Industry benchmark results<sup>1</sup> and the SAS Analytics workload test reflect the scalability and performance of Red Hat Enterprise Linux on the new eight-core processors from Intel. Red Hat and Intel have a long history of collaboration that includes aligning technology roadmaps and making complementary open source contributions. Both companies share a commitment to extending the x86 platform to its full potential.

- **Intelligent performance.** Red Hat Enterprise Linux takes advantage of Intel Xeon processors to adapt throughput to the workload for greater performance. Efficient use of software threads, support for Intel Hyper-Threading Technology, and the ability to change the clock speed on a running processor increases performance as proven by the SAS workload tests.

IN COLLABORATION WITH:



<sup>1</sup> [press.redhat.com/2009/04/16/performance-benchmarks/](http://press.redhat.com/2009/04/16/performance-benchmarks/)



- **Virtualization performance and efficiency.** Taking advantage of next-generation Intel Virtualization Technology, Red Hat Enterprise Linux delivers high consolidation ratios and virtualized performance.
- **Automated energy efficiency.** Red Hat Enterprise Linux uses Intel Intelligent Power Technology to lower power consumption during off-peak times. Red Hat Enterprise Linux includes optimizations designed to operate equipment in the lowest possible power states and throttle processors. This includes exploiting the deep C states available on the Intel Xeon processor 7500 series. Consuming less power means lower cooling requirements, which contributes to further savings and greener datacenters.

**Consolidate with virtualization.** Virtualization enables you to consolidate many SAS application instances onto one server and share resources, resulting in better utilization and performance. Virtualization can also help you migrate to 64-bit or upgrade to SAS 9.2 without interrupting operations.

Red Hat Enterprise Linux Advanced Platform is a leading open source operating platform that delivers built-in virtualization and clustering technologies. It provides advanced capabilities for hosts and guests, including high availability, failover, live migration, storage management, system scheduler, global file system, etc.

- **Increased virtual server scalability.** Support for up to 32 virtual CPUs and 512 GB of memory per virtual machine, and over 400 active workload virtual machines on a 32-core, 1 TB server, enhances flexibility in configuring virtual resources.
- **Open source hypervisor.** Open source contributions from both Intel and Red Hat improve virtualized performance, efficiency, and consolidation ratios, particularly on the Intel Xeon processor 7500 series.
- **Performance.** Up to 140 percent of bare metal performance with multiple virtual machines on a single server compared to a single bare metal operating system. Up to 95 percent of bare metal performance for real-world enterprise applications.

**Lower TCO.** Red Hat Enterprise Linux costs less to acquire when compared to both Microsoft Windows and UNIX. With Red Hat's subscription model, you pay as you go and reduce financial risk with predictable IT costs. Red Hat's value returns to you over the lifetime of the subscription in the form of updated features, security enhancements, and additional hardware and software support. Because you can run and scale even the largest mission-critical applications, such as SAS 9.2 on industry-standard x86 hardware, you reap unprecedented savings on hardware. And more functionality is built into the operating system, reducing the need for expensive add-ons.

#### OTHER RED HAT ENTERPRISE LINUX PERFORMANCE FEATURES:

- Enhancements to the system scheduler provide improved support for shared-cache multi-core systems such as the NEC Express5800/A1080a server.
- Support for CPU sets enables multi-process or threaded applications to be optimally scheduled in large SMP systems.
- Support for a vast application virtual address space on x86-64 systems enables applications such as SAS 9.2 to effectively use more memory per process to work on larger data sets.
- The scalability and performance of large SMP systems is improved to enable applications to more effectively use more processors, reducing resource contention.
- Enhanced I/O schedulers enable I/O performance to be optimized on a per-device basis according to the requirements of the application.
- Support for 10 gigabit Ethernet, iSCSI, and Fibre Channel over Ethernet for high I/O throughput.
- MPIO allows multiple connections from servers to storage to increase availability and throughput. In active/active mode, all paths are used for the I/O, which is spread in a round-robin fashion, increasing the total I/O performance.
- Expanded physical server limits offering support for up to 255 CPUs and 1 TB main memory enhances system scalability.



## MIGRATE FROM RISC TO EIGHT-CORE PERFORMANCE AND SCALABILITY

**NEC Express5800/A1080a server and D3 SAN Storage Array.** Designed specifically for the Intel Xeon processor 7500 series and scalable from 8 to 64 processor cores and up to 128 threads, the NEC Express5800/A1080a server is an ideal platform for scaling up or consolidating SAS 9.2 instances. This server takes advantage of the intelligent performance, energy efficiency, and virtualization capabilities of the Intel Xeon processor and leverages the low latency and high-performance interface technology of NEC's supercomputers. The server supports modular designs, redundant components, hot plug capabilities, and floating I/O. Additional benefits include:

- Up to 64 cores, 128 threads, and 2 TB memory
- Integrated Intel Quick-Path Interconnect technology to increase performance through efficient memory access
- Green Cooling Technology that helps to minimize power consumption and automate power usage for more effective datacenter consolidation benefits
- Built-in service processor working in conjunction with NEC's BIOS and Intel's Machine Check Architecture to provide reliability, availability, and serviceability for mission-critical computing
- Server virtualization that offers high performance, energy efficiency, and higher server bandwidth to handle the increased communications in a virtualized environment

The NEC D3 SAN Storage Array offers scalability to 288 TB and aggregate throughput of over 1,100 MB/s. The system is fully redundant to protect against single point of failure, with a battery backup unit to protect its 4 GB of cache. Features include replication, snapshots, performance monitoring, automatic tuning, multi-pathing, and failover.

**Intel Xeon processor 7500 series.** Built to handle your most processor-intensive, mission-critical applications, the Intel Xeon processor 7500 series delivers a quantum leap

in enterprise computing performance. The Intel Xeon processor 7500 series combines up to eight cores and 16 processing threads in a single device and offers four advanced, high-bandwidth interconnect links that allow multiple processors to be directly connected to each other. The result is unprecedented scalability.

The Intel Xeon processor 7500 series intelligently adjusts performance and energy consumption to accommodate application needs. Built-in Intel Turbo Boost Technology automatically speeds up the processor when your SAS workload requires extra performance. Intel Hyper-Threading Technology allows each processor core to work on two tasks at the same time to enhance performance for highly-threaded workloads. Intel Intelligent Power Technology automatically places CPUs and memory into an optimal power state for maximum performance, while reducing energy use.

### NEXT STEPS

#### [redhat.com/sas](http://redhat.com/sas)

Find the latest information regarding Red Hat Enterprise Linux for SAS solutions, industry events, supporting documentation, and more.

#### [redhat.com/intel](http://redhat.com/intel)

Learn more about how you can benefit when Intel and Red Hat collaborate to offer performance and innovation.

#### [redhat.com/rhel/resource\\_center/reference\\_architecture.html](http://redhat.com/rhel/resource_center/reference_architecture.html)

Explore the Reference Architecture Series to get the most out of Red Hat solutions. Look for a white paper detailing the results and lessons learned during this test, as well as other related testing of SAS 9.2 on Red Hat Enterprise Linux.

#### [redhat.com](http://redhat.com)

Red Hat, the world's leading provider of open source solutions, is headquartered in Raleigh, NC with over 65 offices spanning the globe. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with virtualization, applications, management, and service-oriented architecture (SOA) solutions, including Red Hat Enterprise Virtualization and JBoss Enterprise Middleware. Red Hat also offers support, training, and consulting services to its customers worldwide.

#### [sas@redhat.com](mailto:sas@redhat.com)

Contact us for more information.

### IN COLLABORATION WITH:

