

Red Hat Enterprise Linux: Optimizing Your SAS® Deployments




2012 Orlando Florida

Mike Guerette
Manager
Partner Dev.
Red Hat

Barry Marson
Principle
Perf. Engineer
Red Hat

D. John Shakshober
Director
Perf. Engineering
Red Hat



Agenda

- Who is Red Hat
- Review from last year
- Performance Tuning and Results
 - Intel – High end w/ IBM using SSD
 - IBM – 3 Heterogeneous SAS apps on KVM guests
 - HP – 3 SAS apps in Hybrid config. (2 KVM guests, 1 host)
 - HP – SAS Grid Manager w/ StorNext
- Summary, Q&A, and Learn More



“With RHEL, the O/S is no longer the bottleneck.” - TOM KEEFER, SAS EEC





#1 OPEN SOURCE LEADER

The first to reach \$1B Revenue!

COMPANY REVENUE
FY 2003 - FY 2012

salesforce.com

"The offering of Red Hat Enterprise Linux is fantastic, but what's more important is the company that stands behind it." - PARKER HARRIS, EVP Technology



What We Do

We offer a range of mission-critical software and services covering:



HOW WE DO IT.

We develop everything using an OPEN SOURCE model.

Shared development reduces costs & accelerates innovation.

Open collaboration offers products that genuinely meet customers' requirements.

THE BENEFITS.


- ✓ Flexibility
- ✓ Faster technology innovation
- ✓ Better quality
- ✓ Better price & performance
- ✓ Alignment to your needs

"Red Hat Consulting's local experience backed by Enlogit, a Czech Republic Red Hat Premier Partner, gave us peace of mind during roll out." —PAVEL BARTOS, Lead Analyst, GE Money Bank



Doing FAR more with MUCH less

THE OLD DEVELOPMENT MODEL CAN'T INNOVATE FAST ENOUGH



- 🌐 SAS Data sets are growing exponentially
- 🌐 Timeline to actionable results is shrinking
- 🌐 Demand for deeper insight from analytics
- 🌐 CIO pressure for more use of BI/Analytics

*"The TCO and performance of SAS Analytics on Red Hat Enterprise Linux is a **game changer** that allows our customers to more effectively leverage the power of analytics within the confines of today's shrinking IT budgets."*

—RENEE NOCKER, Director of Technology Product Marketing

Topics from 2011 SGF

SAS on RHEL Customers

- Bank of America
- U.S. National Bank

Technology results – SAS on RHEL

- RHEL 6: 16% performance improvement over RHEL 5
- Results with different file systems: XFS, ext4, ext3, and GFS2
- Performance results: Large SMP (x64) and KVM Virtualization

Separate session from Travelers Insurance

- Now a Red Hat customer reference



“ Performance is easily twice as fast under Red Hat Enterprise Linux. We also reduced our total hardware and operating system licensing costs by 25 percent.”

-Tim Nolan, Manager, Linux Engineering, Travelers



SAS Performance Tuning Guides for RHEL

Understanding I/O Schedulers

- **Deadline (for SAS)**
 - Two queues per device, one for read and one for writes
 - I/Os dispatched based on time spent in queue
- **CFQ**
 - Per process queue
 - Each process queue gets fixed time slice (based on process priority)
- **Noop**
 - FIFO
 - Simple I/O Merging
 - Lowest CPU Cost

Configuring for SAS

- **Ktune service (RHEL5)**
- **Tuned utility (RHEL 6)**
 - throughput-performance
 - latency-performance
 - enterprise-performance (for SAS)
- `# tuned-adm profile enterprise-storage`
- **File system read ahead**
 - Elevate further with **blockdev** command
 - Understand interaction with LVM
- **Transparent Huge Pages (RHEL 6)**
 - Disable where SAS applications run

For more information:

http://support.sas.com/resources/papers/tnote/tnote_performance.html



2012 Orlando Florida


Maximizing SAS Performance with
Intel



#SASGF12

Intel® Xeon® 4S to 8S Performance Scalability

Performance Scalability from IBM 4S to 8S Intel Xeon X7560
Platforms with SAS Mixed Analytic Workload



Intel® Xeon® 4S to 8S Performance Scalability

RUN	Number of SAS Jobs	Run Time	Performance (Jobs/Hour)
4 Socket (32 Cores)	228	42:52	319
8 Socket (64 Cores)	456	42:55	637
Scalability	2.0	1.0	2.0

**Twice more load, twice more resources, constant Run time
= linear scalability!**

RUN	AVG SUSTAINED MB/s	PEAK MB/s	AVG SUSTAINED MB/s per CORE
4 Socket (32 Cores)	2,546	3,992	79.6
8 Socket (64 Cores)	4,527	7,179	70.7
Scalability	1.8	1.8	0.89

- 2.00X performance Scalability from IBM 4S to 8S platform => Linear scalability
- ~2.5 GB/s sustained IO throughput for 4 sockets, ~4.5 GB/s sustained IO throughput for 8 sockets.



2012 Orlando Florida

SAS on Virtual Machines with IBM



#SASGF12

SAS on Virtual Machines

IBM x3850 X5 Server is partitioned into separate Kernel Virtualized Machines:

- **SAS Mixed Analytic Workload** 20 User scenario runs in a 16 logical CPU virtual guest
- **SAS 9.2 OLAP** Server runs in a 16 logical CPU virtual guest serving 50 client connections
- **SPDS** 4.52 supporting 50 clients executes in a 16 logical CPU virtual guest



Virtualized Heterogeneous Workloads

KVM exceptional performance w/ 3 different workloads

1. Mixed Analytic peak I/O rates between 800MB/s and 1GB/s.
2. Over 3500 OLAP queries ran with district manager and unit level queries response times averaging 4-6 seconds.
3. SPDS scenario ran over 1500 queries
 - customer detail query response times averaged less than 1 second
 - unit level query response time were between 4-12 seconds.

Whitepaper available from SAS / IBM / RH upon request.



2012 Orlando Florida

SAS on Virtual Machines at HP Hybrid Configuration



#SASGF12

Problem Statement and Methodology

Problem Statement:

- Can RHEL6 with KVM virtualization simultaneously and effectively run SAS 9.2 Mixed Analytics workload in multiple KVM guests as well as in the host

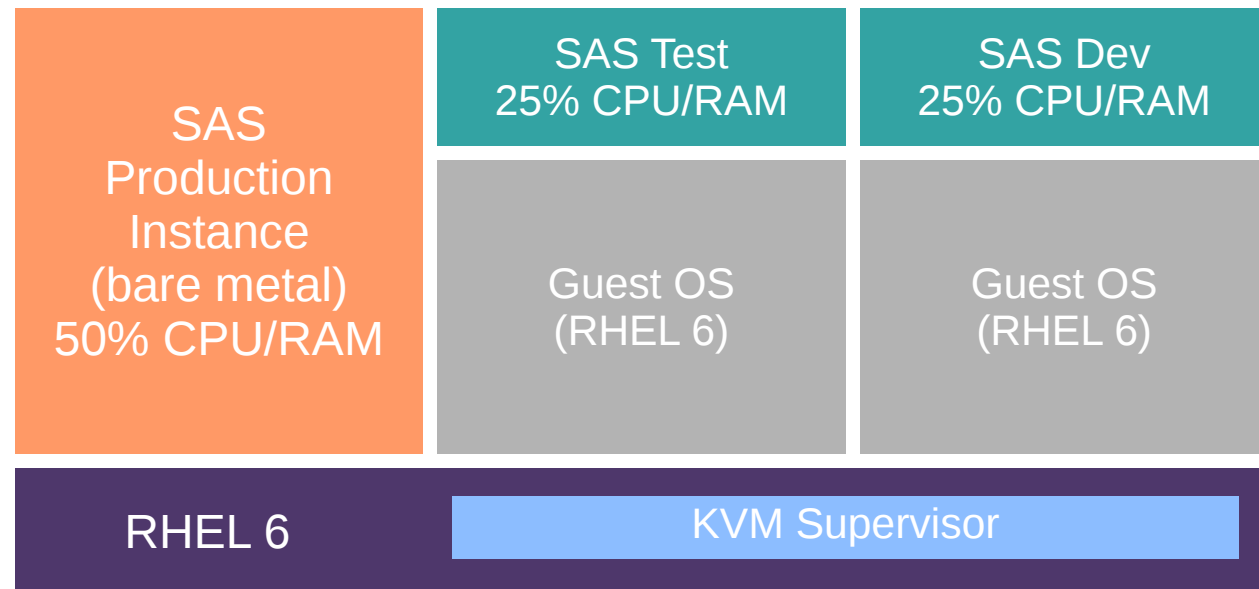
Solution: KVM Hybrid Use Case

Methodology:

- Created a test bed with 3 instances of RHEL 6.1 each running SAS 9.2 Mixed Analytics workload
- All I/O to the guests was through **virtIO** drivers.

Hybrid Use Case


- Unique feature enabled by KVM architecture.
- Mix Bare metal and virtualized workloads on single server
- Isolation, manageability and flexibility
- Save application licensing costs
- Uncompromised performance with agility of virtualization



Results

SAS 9.2 mixed analytics Simultaneous Runs on host/guests Guests are 8 x 64 GB each

HP-DL580g7 - 32 x 256GB - 2 dualport HBA's - 4 HP p2000 arrays
RHEL 6.1 (host/guest)



RHEL Performance Tuning

- Use XFS or EXT4 file systems for best performance
- Set tuned (pronounced Tune-d) to “enterprise-storage” profile
- Read Ahead set to 16K with `blockdev` command
- **HOST**
 - Transparent Huge Pages **enabled** (for guests)
 - `vm.swappiness = 10` (prevents swapping of guests)
- **GUESTS (KVM)**
 - KVM virtIO for each file system. `io=native`, `cache=none`
 - `irqbalance` service disabled (after boot up)
 - Transparent Huge Pages **disabled**

See: Utilize Red Hat Kernel Virtual Machine to Enable Your Development and Test Environment to Peacefully Co-Exist with Your Production Environment

- Bob Augustine, Hewlett-Packard; Barry Marson, Red Hat
- Wednesday, April 25, 09:00 AM - 09:20 AM
- Northern Hemisphere E-3



2012 Orlando Florida

SAS Grid Manager at HP, StorNext, 3PAR



#SASGF12

Benchmark Parameters

Goals

- Achieve maximum performance that an HP Blade environment with 3PAR arrays can deliver
- Show linear scalability based on different SAS Grid workloads
- Use results to come up with Small, Medium and Large reference configurations for HP

Configuration

- Red Hat Enterprise Linux 5.5 installed on 16 blade servers
- HP StorNext File System v4.1.1 shared across 15 blade servers
- HP StorNext Metadata Controller Server - 1 blade server
- SAS Foundation 9.2 and SAS Grid Manager 9.2 installed on 15 blade servers
- Simple, single 30TB file system config – including SAS input, SAS output and SASWork storage



HP SAS Grid Test Highlights – StorNext with 3PAR

15 Blade Server test (LARGE)

- HP StorNext fully utilized the 3PAR T800 – handling 5GB/s on array rated at 5.3GB/s raw
- 144 simultaneous SAS jobs on each Blade Server at same time using SAS Grid team's mixed workload script
- 15 Blade Servers (2x6 core each) running at 25-40% CPU utilization

4 Blade Server test (SMALL)

- 216 simultaneous SAS jobs on each server, 25-30% CPU utilization
- Peak I/O bandwidth of 2.5GB/sec

Details plus reference configurations:

http://h71019.www7.hp.com/ActiveAnswers/cache/78865-0-0-225-121.html?jumpid=reg_R1002_USEN

Looking Forward – Red Hat and SAS

CTO office / Collaboration w/ Performance teams

- Red Hat Enterprise Linux file systems
- Enhancements in EXT3/4 and XFS

Active engineering / evaluation w/ SAS Grid

- GFS2 for use in small scale SAS Grid
- Red Hat Storage (RHS) w/ Gluster

Integration of new OS features

- Transparent huge pages in RHEL6
- Virtualization – VMware / KVM

Summary

Red Hat / SAS Takeaways

- A large and growing percentage of SAS PoC's are on RHEL
- Out-of-the-box Performance with RHEL
- Exceptional Price/Performance for SAS deployments
- Follow SAS Deployment best practices
- Focus on Storage performance first

Next Steps

- Feel confident deploying SAS on RHEL
- Speak to your SAS or Red Hat contacts, and visit our websites.
- Ask for performance briefs, reference architectures, etc.

Visit us at our booth!

Other Global Forum Sessions

- *Storage 101: Understanding Storage for SAS® Applications*
 - Bob Augustine, Hewlett-Packard
 - Monday, April 23, 05:00 PM - 05:20 PM
 - Northern Hemisphere E-3
- *SUPER DEMO: Performance and Scalability with Red Hat Linux and GFS2 Clustered File System in a SAS® Grid Deployment: An Internal Case Study*
 - Ken Gahagan, SAS
 - Tuesday, April 24, 02:00 PM - 02:15 PM
 - Demo Room
- *Utilize Red Hat Kernel Virtual Machine to Enable Your Development and Test Environment to Peacefully Co-Exist with Your Production Environment*
 - Bob Augustine, Hewlett-Packard; Barry Marson, Red Hat
 - Wednesday, April 25, 09:00 AM - 09:20 AM
 - Northern Hemisphere E-3



Key Resources

- Red Hat and SAS: <http://www.redhat.com/SAS>
 - Travelers Insurance – customer case study
 - SAS UNIX-to-RHEL Migration Guide
 - more...
- Red Hat Enterprise Linux:
 - <http://www.redhat.com/rhel>
- Red Hat Reference Architectures:
 - http://www.redhat.com/rhel/resource_center/reference_architecture.html
- UNIX to RHEL Migration:
 - <http://www.redhat.com/sundown>
 - <http://www.redhat.com/rhel/migrate/linux>
- Other Customer References:
 - <http://www.sas.com/success/bofagrid.html>



2012 Orlando Florida

Questions?

Thank You!



#SASGF12