Red Hat Performance Briefs

Red Hat Enterprise Linux 6 KVM and tuned-adm Analysis on Sybase IQ 15.4

Dave Wilson, Sofware Engineer

Version 1.0

June 2012

redhat.

1801 Varsity Drive™
Raleigh NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park NC 27709 USA

The GPG fingerprint of the security@redhat.com key is:
CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

Send feedback to refarch-feedback@redhat.com

Table of Contents

1 Executive Summary

This paper examines the performance characteristics of tuned-adm profiles and Kernel-based Virtual Machine (KVM) on Red Hat Enterprise Linux 6 with Fusion-io technology.

The benchmark utilized represents a Decision Support System (DSS) environment on Sybase IQ 15.4. It's comprised of a series of business-orientated ad-hoc queries with batch updates. The Performance metrics reported are a *query per hour* test that measures the raw query execution power of the System Under Test (SUT) and a *query throughput* test that measures the the ability of the system to process multiple query streams in parallel.

The I/O configuration utilized raw devices, which is recommended by Sybase for IQ, on Fusion ioDrives. These drives were employed for their high throughput and low latency characteristics. For I/O intensive queries these cards can deliver up to 1.5GB/sec of bandwidth for reads and writes.

Red Hat partnered with Fusion-io and Hewlett-Packard for this effort. Fusion-io provided the storage hardware required for the testing and HP provided the server infrastructure.

2 Test Configuration

Commercially available, industry-standard hardware and software was used for the System Under Test.

2.1 Hardware

Server	HP Proliant DL380 G7 2 Socket – 6 Cores (with Hyperthreads) Intel Xeon X5660 @ 2.80 GHz 64 GB RAM (32 GB per NUMA node)
Fusion ioMemory	2 – Fusion ioDrive Duo 320GB

Table 1: Hardware Configuration

2.2 Software

Operating System	Red Hat Enterprise Linux 6.2 (kernel-2.6.32-220.el6.x86_64)
Database	Sybase IQ 15.4
Fusion ioMemory	Firmware v5.0.6 rev 101583 Fusion ioMemory driver version: 2.3.0 build 281

Table 2: Software Configuration

3 Testing - Phase 1

The goal of Phase 1 was to determine the performance effects of various tuned-adm profiles on a DSS type workload.

The tuned-adm tool allows users to easily swap between a number of profiles that have been designed to enhance performance for specific use cases. The profiles that are particularly useful in improving storage performance are latency-performance, throughput-performance and enterprise-storage. Table 1 details profile attributes.

tunable	default	latency-performance	throughput-performance	enterprise-storage
I/O elevator	CFQ	deadline	deadline	deadline
CPU governor	OnDemand	performance	performance	performance
kernel.sched_min_granularity_ns	4ms	4ms	10ms	10ms
kernel.sched_wakeup_granularity_ns	4ms	4ms	15ms	15ms
Disk read-ahead	1x	1x	4x	4x
vm.dirty_ratio	20%	20%	40%	40%
Filesystem barrier	on	on	on	off

Table 1

3.1 Tuned-adm Profile Testing Procedure

The following procedures were applied to Phase 1 testing:

- The Sybase IQ database was configured to a 300GB data set size.
- Install tuned and tuned-adm
  - # yum -y install tuned
- Enable tuned
  - # chkconfig tuned on
  - # service tuned start
- Select profile to run
  - # tuned-adm profile enterprise-storage
- An individual stream of 22 queries were run in a serial with updates done at the beginning of the stream and at the end of the stream to simulate a batch update environment. Query and update times were utilized in an algorithm to derive the query per hour metric.
- A series of 5 streams of 22 queries were run in parallel with the longest running query being used to compute the query throughput metric. Batch update streams were also done in a serial mode during run.

3.2 Tuned-adm Results

Figure 1 details results between tested profiles and the default setting.

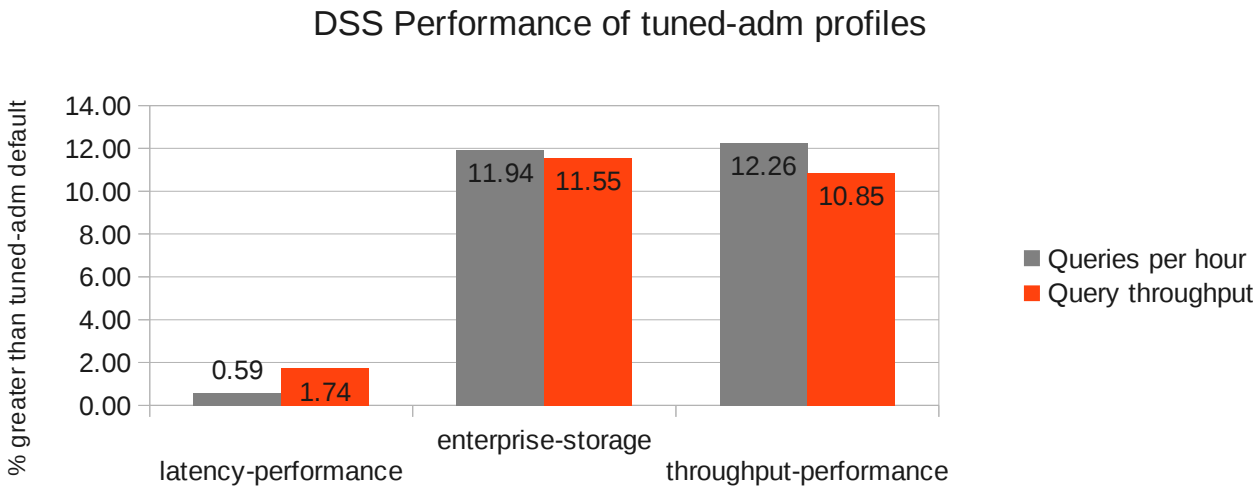

Figure 1

3.3 What does it all mean?

As detailed in the above graph latency-performance had minimal performance advantage over the default settings. Enterprise-storage and throughput-performance profiles had a noticeable gain when deployed.

It should be noted that fusion ioDrive technology utilizes the noop I/O elevator and since raw devices are used, file system barrier, disk-readahead and VM dirty ratio settings have no effect on the measurements. However, the profiles would still be recommended for applications that are I/O intensive and reside on traditional fibre channel and file systems.

The two profile attributes that do have an impact on the testing are related to the Linux completely fair scheduler (CFS):

sched_min_granularity_ns - specifies the minimal preemption granularity for CPU bound tasks.

sched_wakeup_granularity_ns - Increasing this variable reduces wake-up preemption, reducing disturbance of compute bound tasks. Which could assist in some of the CPU bound queries.

4 Testing – Phase 2

The goal of Phase 2 was to determine the virtualization overhead of KVM on the DSS workload with single and multiple guest environments. Virtualization results are then compared with bare metal results.

4.1 KVM configuration and testing procedure

The following procedures were applied to Phase 1 testing:

- The Sybase IQ database was configured to a 300GB for a single VM (virtual machine) test and 150GB for a 2 VM test data set size.
- The memory and allocated cpus were halved for the 2 guest run (12 vcpu's and 32GB RAM per guest).
- kvm tunables used were virtio drivers for disk, cache=none, aio=native, and type=raw.
- Numa tuning was achieved using libvirt parameters cpuset and numatune on the 2 guest run.
- An individual stream of 22 queries were run in a serial with updates done at the beginning of the stream and at the end of the stream to simulate a batch update environment. Query and update times were utilized in an algorithm to derive the query per hour metric.
- A series of 5 streams of 22 queries were run in parallel with the longest running query being used to compute the query throughput metric. Batch update streams were done in a serial fashion during run.

4.2 KVM Results

Figure 2 details results between tested profiles and the default settings.

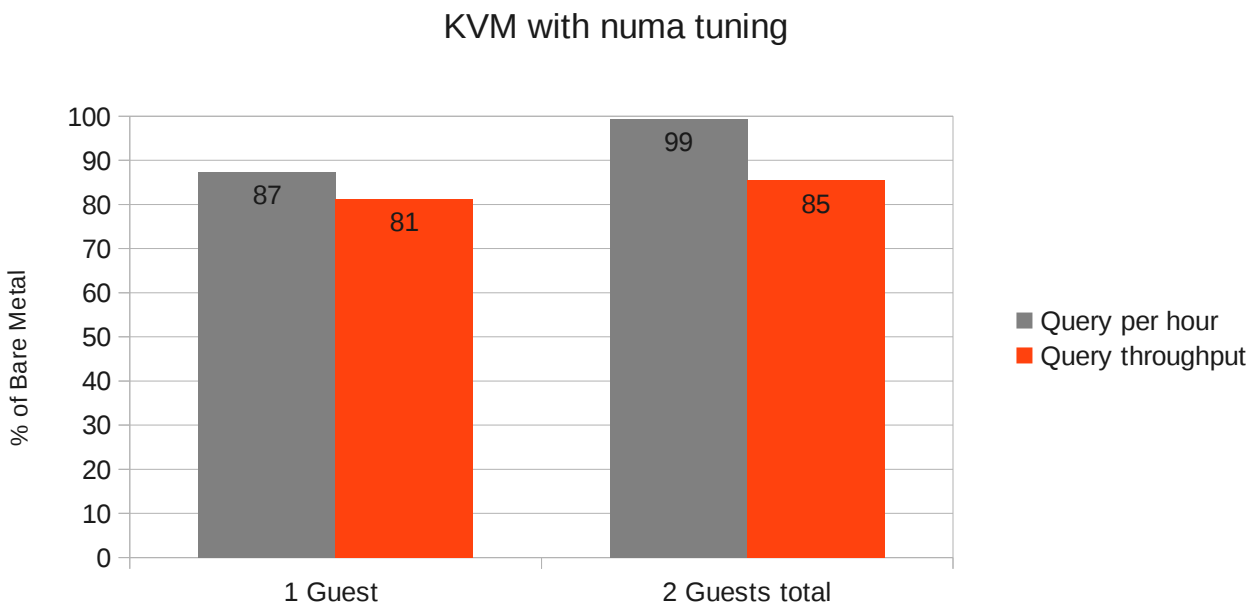

Figure 2

4.3 What does it all mean?

With KVM and Red Hat Enterprise Linux 6.2 you can deploy applications in a virtualized environment and increase server utilization without sacrificing performance. As the above data confirms, virtualized guests can achieve over 80% of bare metal results with a single guest on the query per hour and query throughput tests. On a two guest environment, the sum of two guests achieved 99% of bare metal with query per hour and 85% of query throughput tests.

5 Conclusion

Data collected during the tuned-adm testing phase demonstrate how a properly selected tuned-adm profile can significantly improve application performance with Red Hat Enterprise Linux 6.2. Phase 2 testing shows that server consolidation with KVM can be achieved without sacrificing performance, providing close to bare metal performance on Enterprise Class applications.

Appendix A: Revision History

Revision 1.0	Tuesday June 05, 2012	Dave Wilson
Initial Release		