KNOWLEDGE consulting group KCG

11710 Plaza America Dr., Ste 520, Reston, VA 20190
Voice: 703.467.2000 Fax: 703.547.0322
www.KnowledgeCG.com

CSN16: ArcSight ESM Performance Tuning with RHEL

Joe Burke
Solutions Architect

Overview

- How to design and tune your ArcSight ESM system using Red Hat Enterprise Linux (RHEL) to achieve maximum performance.
- How to properly create a test plan to accurately measure the performance impact of any configuration change.
- How to simulate user and event source activity in a steady and predictable manner in order to ensure that the metrics gathered are reliable enough to support configuration changes and hardware upgrade decisions.
- Review specific tuning examples to improve report, active channel, and trend performance.

The value of tuning

- Better ArcSight ESM performance increases productivity
  - Less time spent waiting for results
- Increased return on your hardware investment
  - The 4 step tuning example below shows the tuned configuration returning query results twice as fast as the default install
- Similar to tuning a race car
  - Changes that have a positive impact are added to the baseline
  - The larger changes should be tested before smaller changes
    - IT example: If you plan to test different storage arrays and also different host bus adapters, test the different storage arrays first, then proceed to test host bus adapters using the best performing storage array.
    - Race car example: Test the engine with different pistons first, then test different spark gaps using the best performing piston.

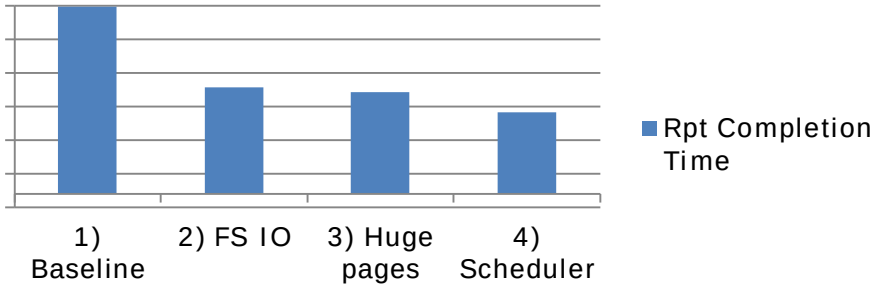

The scientific method

1. Ask a question
   A. This question is usually “How can I improve ArcSight ESM performance?”
2. Do background research
   A. Find information about things that might improve ArcSight ESM performance.
3. Construct a hypothesis
   A. Based on background research and observation of the system, make an educated guess that changing one specific part (a variable) of the system will positively affect ArcSight ESM performance.
4. Test your hypothesis by doing an experiment.
   A. It is wise to run the experiment at least 3 times to ensure the first results weren’t an accident.
   B. Verify that none of the parameters outside of what you are testing have changed between test runs(like shmmax)
5. Analyze your data and draw a conclusion
   A. Collect your data and analyze it to see if your hypothesis was true or false.
   B. If it was false, construct a new hypothesis and try again.
6. Communicate your results
   A. Share with the ArcSight community!
   B. https://protect724.arcsight.com

Background research

- With the diversity of RHEL and Oracle performance recommendations available, which recommendations specifically help ArcSight ESM?
- Here are two good examples to start with
  - Oracle 10g Server on Red Hat Enterprise Linux 5 Deployment Recommendations http://www.redhat.com/f/pdf/rhel/Oracle-10-g-recommendations-v1_2.pdf
  - ArcSight ESM Performance Guide: https://protect724.arcsight.com/docs/DOC-1198
- Cross-reference the parameters you're researching whenever possible
  - For example, using filesystemio_options=setall is mentioned in the ArcSight ESM 5 server.defaults.properties and the Oracle 10g Server on Red Hat Enterprise Linux 5 Deployment Recommendations guide.
- Testing and tuning allows you to identify what works best in your environment

General guidelines for setting up a test environment

- Remove software that may interfere with testing
  - Antivirus
  - Host intrusion detection/prevention
  - Firewalls
  - Non-essential logging
- Disable any scheduled tasks or cron jobs that could impact your test results
- All system conditions must be as stable as possible

Initial installation of ArcSight ESM

- Recommend starting with the XLarge template as a base
- Test environments often do not have as much storage capacity as production, therefore you may be forced to use an extremely short retention period. I found the following settings work well, especially if you're experimenting with extremely high event rates that require shorter retention periods to avoid running out of free space.
  - Online retention = 3 days
  - Online Reserve = 7 days
  - Compression waiting period = 2 days
- Disable default content
  - When asked for additional packages, always choose "none"
  - Disable "ASM Database Free Space" and "ArcSight User Login Trends – Hourly" trend
  - Disable all threat tracking rules
  - Disable "ArcSight User Sessions", "Current Users Logged In", and "User Access Log" Data Monitors
- Do not enter a URL for ArcSight Web

Modifying database templates

- The default location for the XLarge template is:
  - /usr/local/arcsight/db/installer/oracle10g/unix/dbca/ArcSight_XLarge.dbt
- You can modify the memory parameters depending on the amount of memory installed in your DB server.

<DatabaseTemplate name="ArcSight X Large Template" description="

Total Available Memory: 8G

Total Required Memory: 6096M

<initParam name="memory_target" value="6096" unit="MB"/>

- This is your opportunity to change the db_block_size if you wish.

Reconfiguring ArcSight ESM between tests

1. Create an ARB of the content you want to re-use for future tests if necessary.
2. Rundatabasesetup
   A. Initialize ArcSight Tablespaces, Schema and Resources
   B. Create/Recreate Schema Only
3. arcsight managersetup
   A. Follow the prompts and use the same parameters as the last test run
   B. Most prompts will already be filled in for you.
4. Disable default content
   A. When asked for additional packages, always choose “none”
   B. Disable “ASM Database Free Space” and “ArcSight User Login Trends – Hourly” trend
   C. Disable all threat tracking rules
   D. Disable “ArcSight User Sessions”, “Current Users Logged In”, and “User Access Log” Data Monitors
5. Add your ARB to install scheduled reports or other testable content
6. Re-register the test alert connector, and start the service at the specified time the test should begin.
7. Record results at the end of the test and return to step 1.

Testing content and simulating users

- Scheduled reports
  - Offer the ability to execute as many simultaneous queries as you need
  - The report, query, filter, even the schedule can be saved in an ARB to re-use over and over again
  - Makes it easy to ensure that everything is run exactly the same way for each test

- A production system is always doing multiple things, so it's a good idea to put your test system under additional user load beyond the specific focus of your testing.
  - The additional load should be consistent and predictable. For instance, I usually have a few different reports run several times per hour every hour.

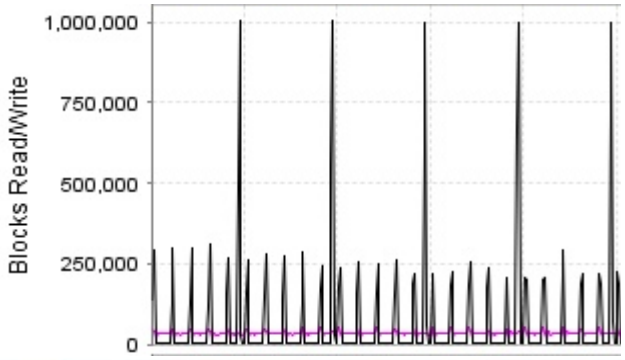

Using the Test Alert connector

- Read the “ArcSight Event Replay How-To Guide” for details on the Test Alert connector:
  - https://protect724.arcsight.com/message/16799
- This document explains how to capture events from the devices relevant to your environment and replay them at a controlled rate to your test system.
- If you’d prefer to use “Bleep”, instructions for that tool can be found here:
  - https://protect724.arcsight.com/docs/DOC-1198

Ways to monitor performance

- Sysstat and kSar

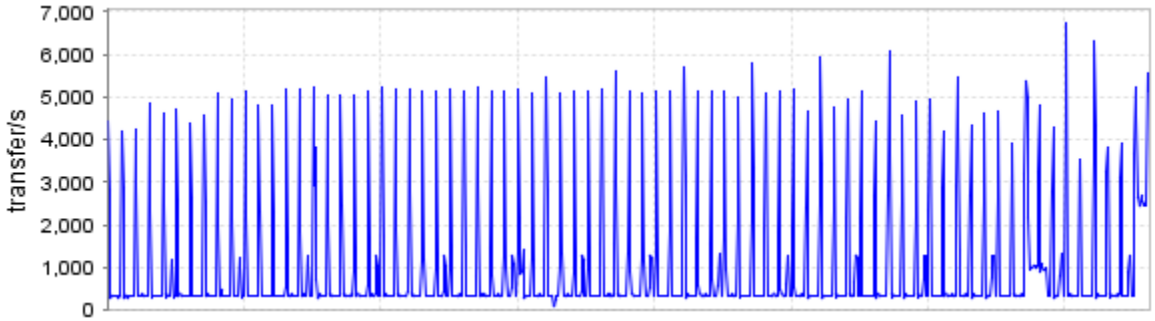

- Internal ESM events

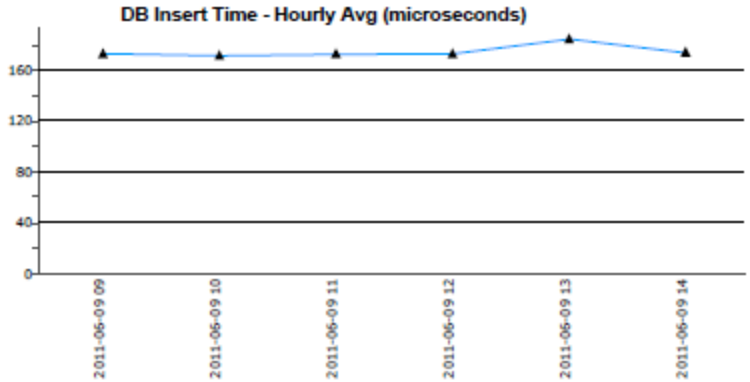

Sysstat

- Before you start using kSar you must first install the sysstat RPM (which is listed as a required package in the ArcSight ESM installation guide, so you probably already have it) on your database server which will then automatically gather system statistics for you.

- Disk statistics are not collected by default, but it's easy to enable it. To enable disk statistics collection:
- Just add a "-d" in two files, then restart sysstat

- vi /etc/init.d/sysstat
- /usr/lib64/sa/sadc –F –L –d - && touch /tmp/sysstat.run

- vi /etc/cron.d/sysstat
- */10 * * * * root /usr/lib64/sa/sa1 –d 1 1

kSar

- An open source tool available from: http://sourceforge.net/projects/ksar/
- Pull stats directly from a server and display them graphically
  - sar -A -f /var/log/sa/sa14 -s 10:30:00 -e 13:30:00
- Pull stats from a txt file
- Select specific time ranges to graph
- Export data as text or as an image
- To get the actual disk device name, you'll need to do "cat /proc/partitions". Use "Options -> disk name" to assign an easy to recognize alias to any of the disks.

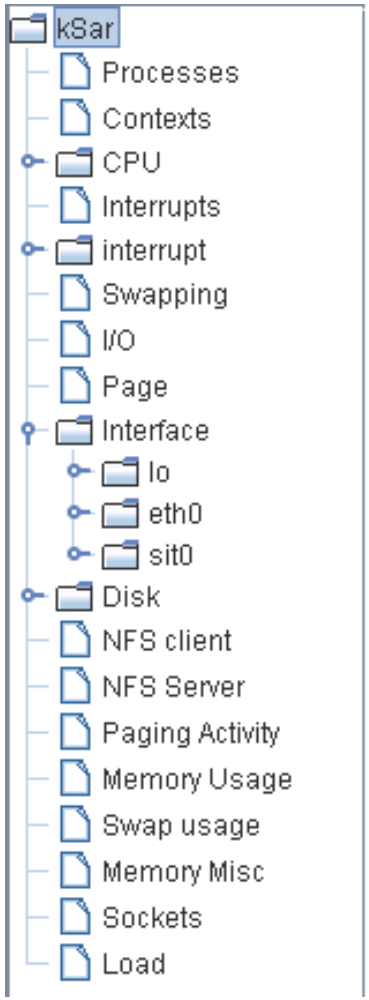

Disk activity – Interesting observations

- From the Oracle White Paper: Optimizing and Protecting Storage with Oracle Database 11g Release 2

"Size disk arrays according to the workload they need to support

Traditionally, storage arrays used underneath Oracle Databases have been sized based on the amount of data that needs to be stored in corresponding databases. As disk capacity has increased, and processor cores have become faster, this sizing metric is no longer useful. Instead, storage arrays should be size based on the performance they can deliver to database applications. Oracle Database workloads consist of random I/O (input/output) operations, that are typically caused by data manipulation operations such as inserts, updates and deletes; and sequential I/O operations that are caused by queries. These operations are typically measured as follows:

-Random I/Os are measured in I/O Operations per Second, or IOPs.

-Sequential I/O is measured in the number of megabytes of data that can be scanned from the storage system per second, or MB/s.

Storage Arrays should be sized on the number of IOPs they deliver, as well as the achievable MB/s that they can deliver."

- Since IOPS and MB/s are important metrics. How can they be observed and measured?

REDO volume (Contains only the Oracle redo logs)

- High write activity with almost zero reads. (When archiving is disabled)

UNDO volume (Contains only the ARC_UNDO datafiles)

- High write activity with peaks appearing during a query.
- Some read activity
- Write activity at times will exceed REDO

DATA volume (Contains only the ARC_EVENT_DATA datafiles)

- Large spikes of read activity during queries that use unindexed fields in their conditions.
- Some write activity, but far less than REDO or UNDO.

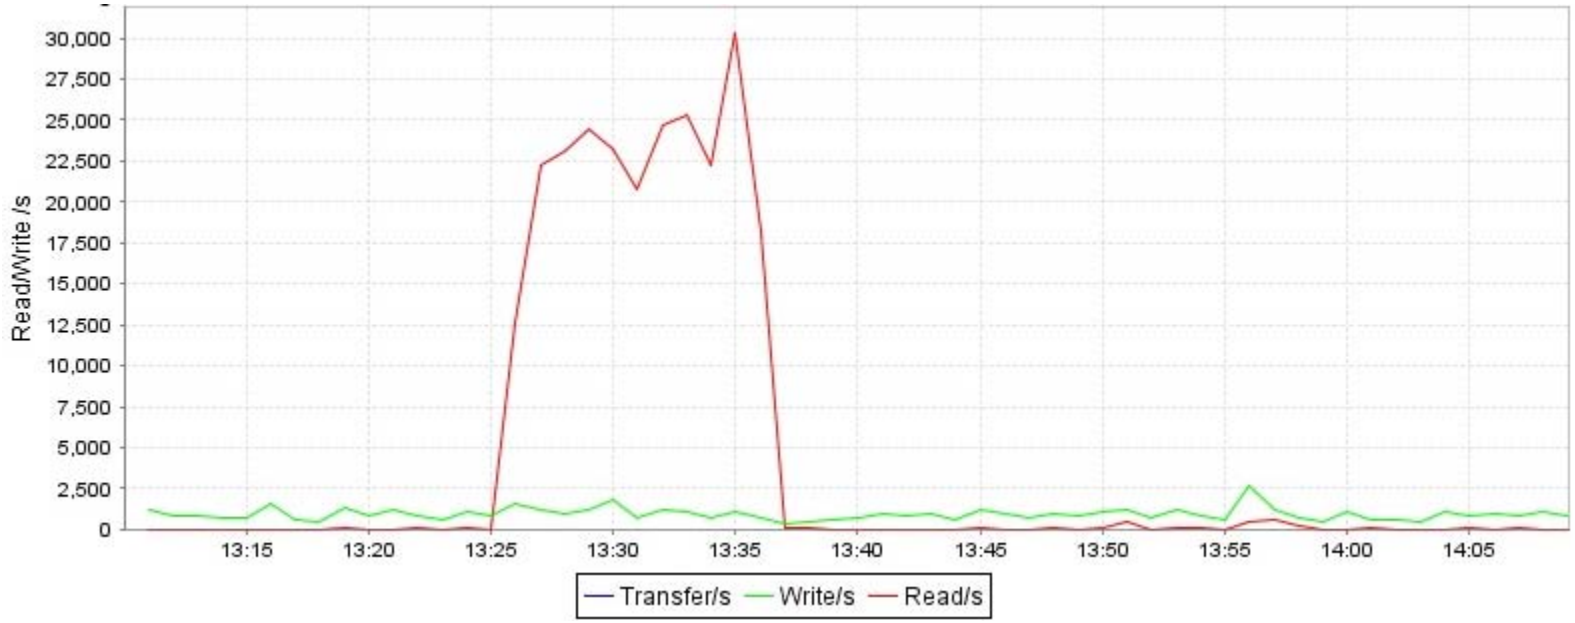

INDEX volume (Contains only the ARC_EVENT_INDEX datafiles)

- HUGE spikes of read activity during queries that use indexed fields in their conditions.
- Some write activity, but far less than REDO or UNDO.

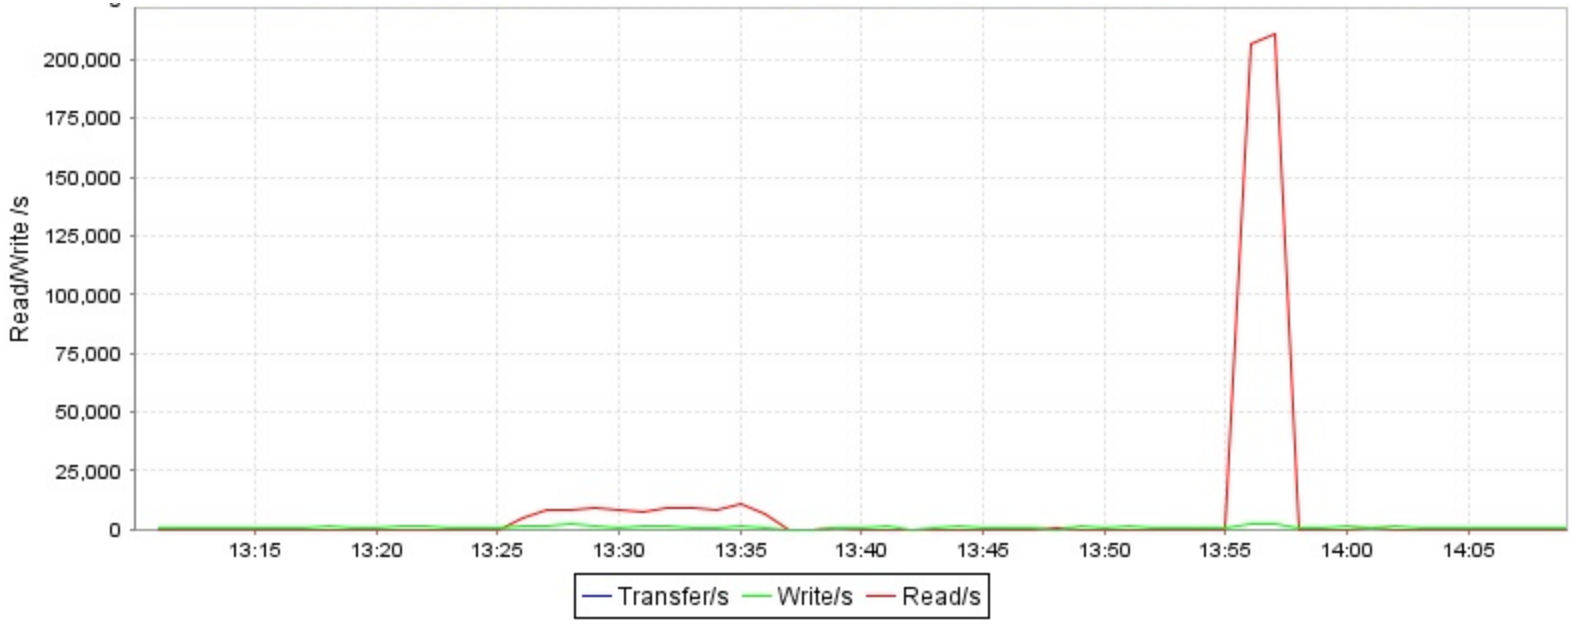

The “everything else” volume

- This graph shows activity for the volume holding all remaining files for the ArcSight ESM DB.
- This includes the RHEL OS, Oracle home directory, all other Oracle datafiles, etc.
- The activity is extremely low compared to the other volumes

Oracle datafile placement

- Oracle only cares about the name of the mount point.
  - You can stop the oracle service, copy the datafiles to a temporary location, rebuild the partition using more disks, and then copy the datafiles back.
- You can also move individual datafiles using a combination of Oracle SQL and RHEL shell commands
  - 1. Login to SQLPlus.
  - 2. Connect as SYS DBA with CONNECT / AS SYSDBA command.
  - 3. Shutdown the database instance with SHUTDOWN command.
  - 4. Rename or/and move the datafiles at operating system level.
  - 5. Start Oracle database in mount state with STARTUP MOUNT command.
  - 6. Modify the name or location of datafiles in Oracle data dictionary using following command syntax:
  - ALTER DATABASE RENAME FILE '<fully qualified path to original data file name>' TO '<new or original fully qualified path to new or original data file name>';
  - 7. Open Oracle database instance completely with ALTER DATABASE OPEN command.

Different query, different datafiles

- These graphs show a query with an unindexed field at 13:25, and an indexed field at 13:55.

Index Volume

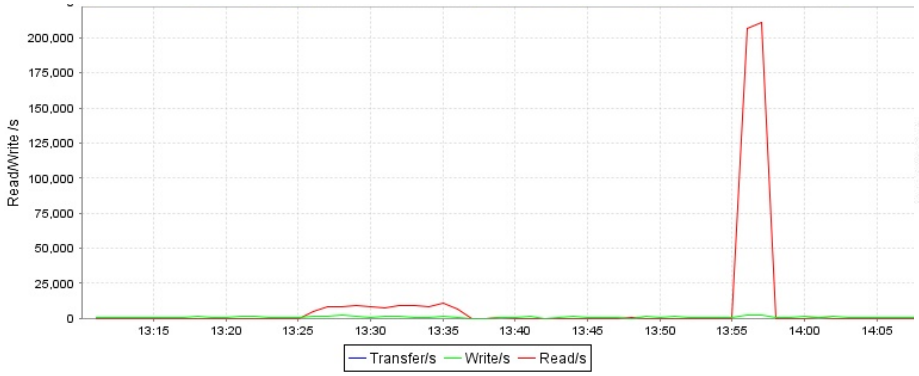

Data Volume

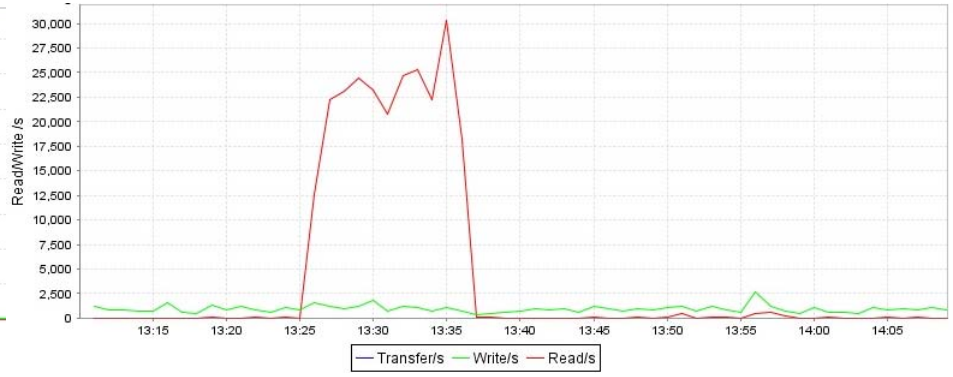

- One possible strategy to optimize performance based on this information would be to index more fields and place the index datafiles on faster storage.

ORION

- The hypothesis is that “increasing IOPS will decrease query time”
- Background research
  - http://www.dell.com/downloads/global/products/pvaul/en/deploying-oracle-ssd.pdf
- ./orion_linux_x86-64 -run advanced -testname orion -write 65 -size_small 32 -cache_size 4096 –duration 120
- You’ll find the output from ORION as a file ending with “summary.txt”
  - Maximum Large MBPS=75.33 @ Small=0 and Large=2
  - Maximum Small IOPS=348 @ Small=5 and Large=0
  - Minimum Small Latency=5.76 @ Small=1 and Large=0
- Your generic tuning goal is to get the lowest latency, highest IOPS, and highest MBPS with a 32k block.
- Variables include: stripe sizes, LUN sizes, short-stroking, drive type, number of drives, etc.
- When comparing HDD and SSD, you may only see improvements in IOPS and latency yet query performance still improves.
- **WARNING: This example uses the “-write” switch which will erase the target volume during testing**

ORION sample results

- This volume holds data, index, and undo tablespaces
- Lab Config 1 – (2 disks, RAID0)
  - Maximum Large MBPS=40.18 @ Small=0 and Large=2
  - Maximum Small IOPS=200 @ Small=5 and Large=0
  - Minimum Small Latency=7.52 @ Small=1 and Large=0
- Lab Config 2 – (3 disks, RAID0)
  - Maximum Large MBPS=45.55 @ Small=0 and Large=2
  - Maximum Small IOPS=292 @ Small=5 and Large=0
  - Minimum Small Latency=6.96 @ Small=1 and Large=0

Putting it all together

200 IOPS (Query completed in 80 seconds)

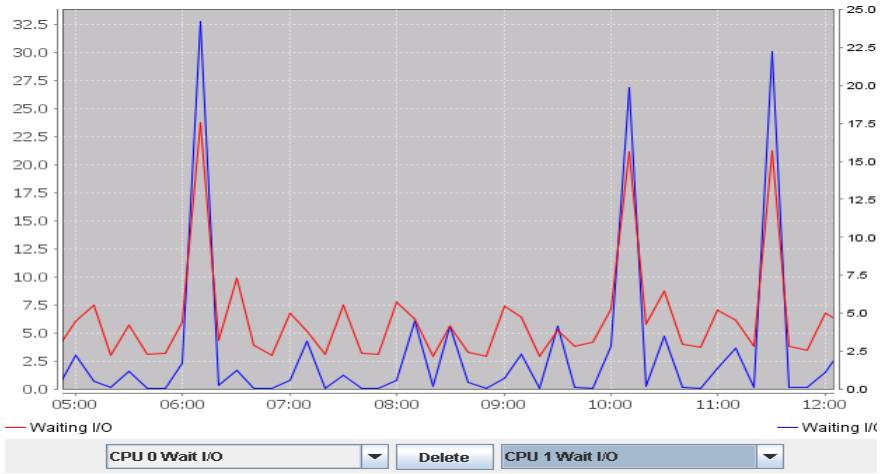

292 IOPS (Report completed in 50 seconds)

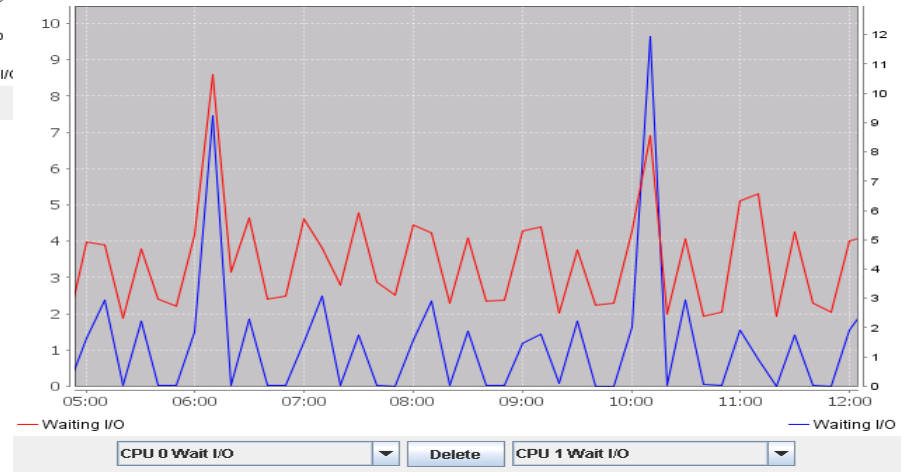

The effects of write activity

- ORION was executed against a volume that contained the redo logs
- ./orion_linux_x86-64 -run advanced -testname orion -size_small 32

- These are the results from ORION without any activity
- Maximum Large MBPS=22.43 @ Small=0 and Large=2
- Maximum Small IOPS=92 @ Small=5 and Large=0
- Minimum Small Latency=14.45 @ Small=1 and Large=0

- These are the results from ORION with 1000 EPS from the Test Alert connector
- Maximum Large MBPS=19.59 @ Small=0 and Large=2
- Maximum Small IOPS=82 @ Small=5 and Large=0
- Minimum Small Latency=16.09 @ Small=1 and Large=0

Configuration parameters known to enhance ArcSight ESM performance

- filesystemio_options=setall (Oracle)
- Huge Pages (RHEL)
- Elevator=deadline (RHEL)

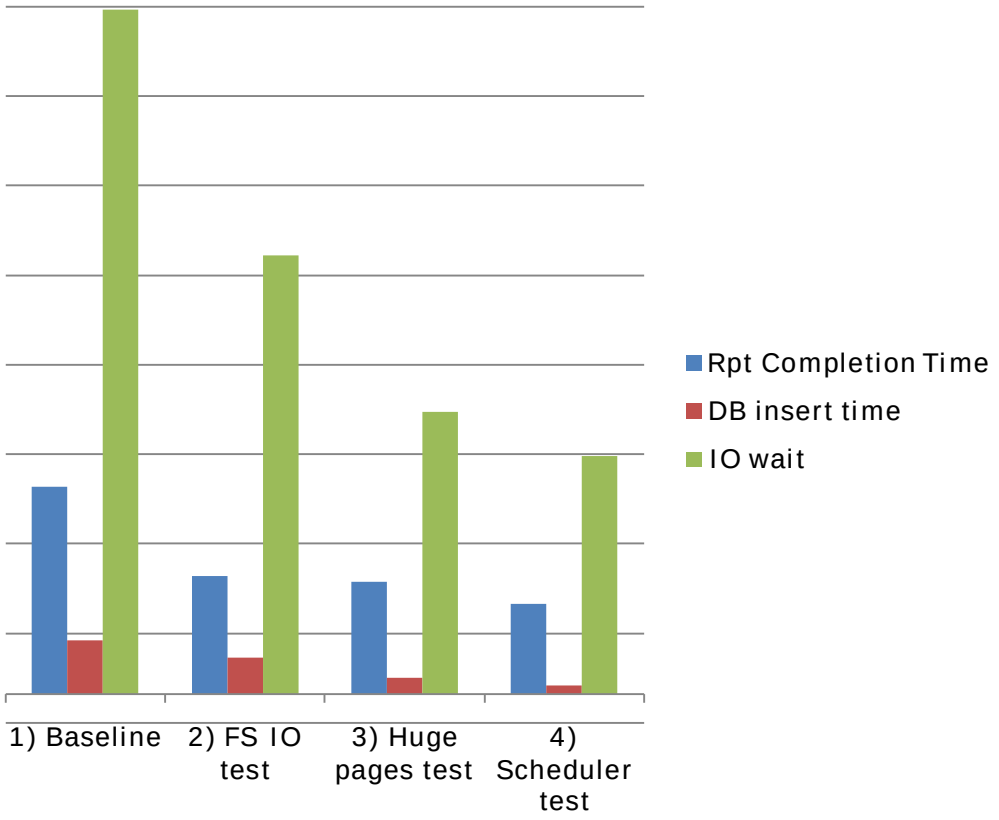

Filesystemio_options

- filesystemio_options=setall with EXT3 offers 97% the performance of raw
  - http://www.redhat.com/magazine/013nov05/features/oracle/
- filesystemio_options=setall is mentioned in the ArcSight ESM 5 server.defaults.properties and the Oracle 10g Server on Red Hat Enterprise Linux 5 Deployment Recommendations guide.

Swapping

- Swappiness problem fixed by setall (swap usage is ok, frequent swapin/swapout is bad)

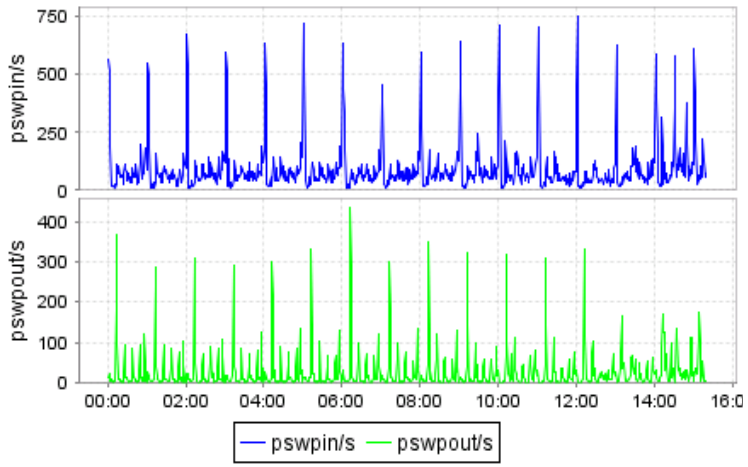

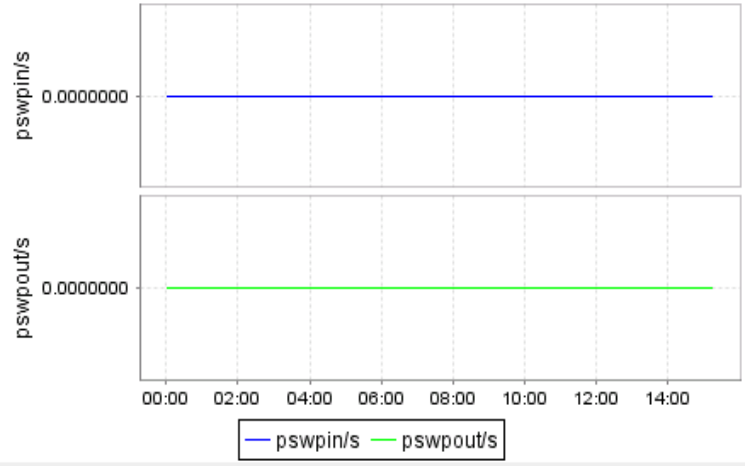

- This is an example of why you should test the larger change (setall) first before testing smaller changes like swappiness.

Huge Pages

- Recommended by Red Hat
- AMM and ASMM must be disabled in order for huge pages to work
- sga_max_size: Choose this value first based on the amount of memory you want Oracle to use
- /etc/sysctl.conf
  - kernel.shmmax: A value in bytes, slightly higher than sga_max_size
  - kernel.shmall: shmmax / 4096
  - vm.nr_hugepages: shmmax / 2048 / 1024
- /etc/security/limits.conf
  - oracle soft memlock: shmmax/1024
  - oracle hard memlock: shmmax/1024

The scheduler

- /boot/grub/grub.conf
- Add “elevator=deadline”
- Recommended by Red Hat

DBMS Stat Tuning

- The value assigned to these parameters can have a profound impact on performance.
- select * from sys.aux_stats$ where sname='SYSSTATS_MAIN';
- GatherSystemStats.sql (automatically set stats using the ArcSight supplied script)
  - EXEC DBMS_STATS.GATHER_SYSTEM_STATS;
- Setting stats manually using SQL
  - exec dbms_stats.set_system_stats('CPUSPEEDNW',1887);
  - exec dbms_stats.set_system_stats('IOSEEKTIM',10);
  - exec dbms_stats.set_system_stats('IOTFRSPEED',4096);
  - commit;
- Read ArcSight KB 3903 for more details

Report tuning using hints

- You can add a hint to your query and then choose that query for your report
  1. No hint (baseline)
  2. Hint = FULL(Event)
  3. Hint = FULL(Event) PARALLEL(AUTO)
- In this example, a 36 hour query using an unindexed field showed big improvements from using a full hint. Adding parallel execution had a negative impact.
- To permit the use of hints in a query, edit \current\config\console.properties and add the following line
  - database.hint.editable=true

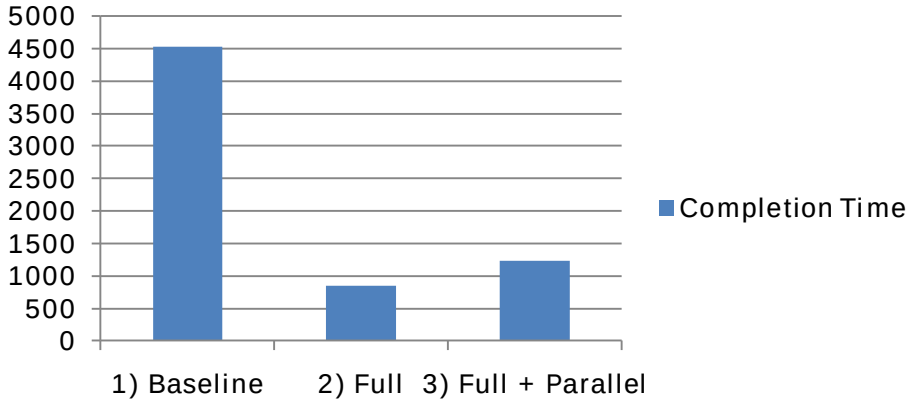

Query performance comparison - hints

- A query using a non-indexed field with a full hint was executed at 7:15. Notice the lack of activity on the volume where the ARC_EVENT_INDEX datafiles reside.
- A query using an indexed field was executed at 9:00. Notice the lack of activity on the volume where the ARC_EVENT_DATA datafiles reside.

Index Volume (Full hint)

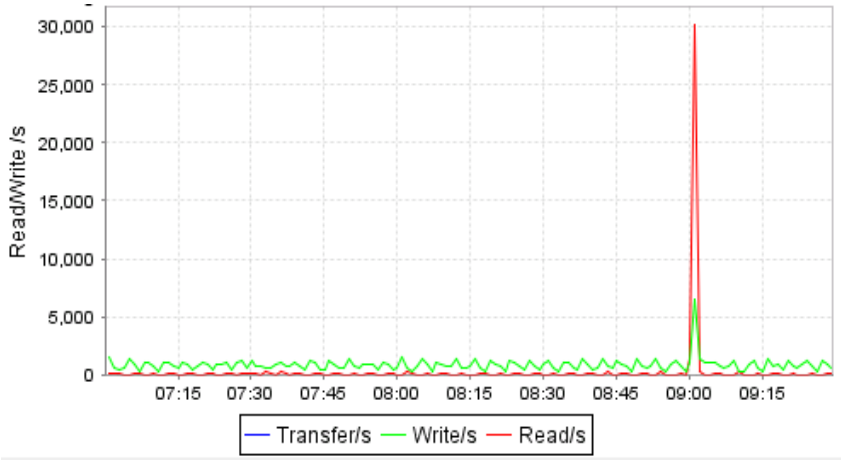

Data Volume (Full hint)

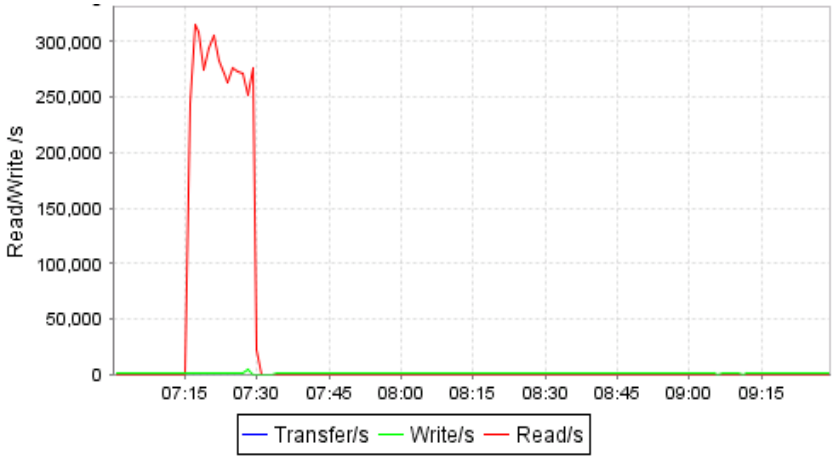

Document your system

- Manager and DB server hardware
- Oracle datafile locations and physical storage layout
  - Drives -> RAID Group -> LUN -> Partition
- Config files
  - /etc/sysctl.conf
  - /etc/security/limits.conf
  - /etc/fstab
- Partition management runtimes
  - Manager
  - Compressor
  - Archiver
- Retention period
  - Online
  - Reserve
  - Compression waiting period
- Oracle parameters
- server.properties

Additional performance tuning ideas

- Hardware upgrades
  - More memory or faster memory (Usually best for DB server)
  - Faster processors (Usually best for Manager server)
  - Vendors may be willing to provide a free evaluation server or storage device
- Storage array changes
  - Additional fibre channel paths
  - Additional controllers
  - Additional drives (this usually has the biggest impact)
  - Try different FC HBA drivers and configuration options
- Research using other sessions
  - https://protect724.arcsight.com/community/protect10/sessions
    - SN62: Gain Rock Star Status as an ArcSight ESM Manager Administrator
    - SN71: ArcSight ESM Database Performance from the Bottom-Up

Tracking/Tuning performance in production

- Very difficult to measure small improvements due to the large number of variables
- It may be possible to get a rough estimate of performance by tracking IOwait for all processors using a daily average.
- Increased IOwait could mean
  - Increased event rates
  - Increased user activity
  - Decreased storage performance
- The best way to improve the performance of a production system is to have an identical test system that you can test your hypothesis with and apply changes that have a known good impact to production.

RHEL 6

- http://www.redhat.com/summit/2011/presentations/bestof/thursday/shak_larry_t_1020_Perf_summit2011.pdf
- ESM 5.0.1 w/ Oracle 11g install appears to work
- Improvements to huge pages
  - NUMA aware Hugepages
  - Transparent hugepages
  - 1GB hugepage support
- 10Gb networking improvements

Questions?