

Q&A

Service Manager Performance Tuning Deep Dive

Questions Asked by Attendee

Answers

Q: Is there an upper limit for RAM that will provide little to no return, e.g. beyond 16GB?	I am not aware of any downside to having too much memory. The only issue would be the extra heat that the memory sticks generate, but modern operating systems do a nice job of using extra memory for caching frequently-used files, so more memory only benefits them.
Q: I'd like to setup IR in the background, and have tried to do this in the past. How is the processing handled, by a schedule record? Do I create that schedule record?	Put <code>ir_asynchronous:1</code> in each <code>sm.ini</code> file of all app servers and <code>sm -que:ir</code> on your server running the load balancer process. All the servers will need to be restarted. There is no schedule record, the <code>ir</code> que process will scan the holding table for records to process automatically.
Q: Do you mean a separate JVM when you are talking about a 'servlet'?	Yes, each servlet is it's own JVM container, including the background processes, which run in their own servlet container also.
Q: Do you think that HP will really stop support of 7.02 in September?	Not sure. If you call HP, they may help you, but you should really upgrade. Perform at least a binary upgrade to at least SM 7.11 patch 13. You do not HAVE to upgrade your RAD apps as they will still run in SM 7.11 (and probably 9.20 and 9.21)
Q: Our intel team patch servers every month, but last month patch did something with memory and SM wouldn't start...lol Had to define a shared memory address for it to start. Q: there is a parameter (<code>ir_disable:1</code>) you can use in <code>sm.ini</code> to disable <code>ir</code> as needed Q: HP Server sizing doc for SC6 and SM.7.0 on windows says 30 users per servlet and 50 for unix	30 users per servlet for 32 bit operating systems may be better than 50. Shouldn't have any problems running 50 in 64 bit OS
Q: Assuming that I have lots of memory, i.e., your recommendations, how would you recommend I set the Java memory parameters for the individual servlets?	Have had good luck with <code>-JVMOption0:-Xms512m -JVMOption1:-Xmx512m</code>
Q: Isn't Service Manager 7.11 defaulted to 256meg or maybe 512meg java heap. How can it use 2 GB per servlet then? I realize other things than just the heap are in the total JVM memory. We have upped our servlets to use a 1 GB heap and HP is threatening to not support it anymore (to high, they say go back to 256meg/servlet)	Java memory is very complex and you can find 'short' articles of hundreds of pages of how to set params for garbage collection, etc. Heap size is only one parameter for java memory. Servlets will still use up to 2 gb with the heap size set to 256 or 512 MB
Q: Is there a problem with tables with too many columns? Q: Record lists do a <code>SELECT <columns></code> while Views do a <code>SELECT *</code> resulting in performance differences. Any workaround? Q: Is there a problem with having too many multi row array tables associated to a single dbdict? Q: SQL to SQL mapping often crashes on large tables, any workaround? (e.g. <code>cm3r</code> on 9.20)	1. No. Although there will be more joins, typically there won't be an issue. The only caveat is on tables like <code>inbox</code> where the primary key is not a single field and those fields are also used in the 'where' clause. 2. HP is working on the Views issues. They currently are extremely inefficient 3. No 4. You'll have to identify the root cause. Hopefully it will be in the log. You may have to log the entire sql remap.
Q: we frequently come across the iss with SCEMAIL service was hung and need to restart to sent the mail from <code>sc.6.2</code>	Consider switching to <code>scauto</code> for email that runs as a separate service that you can restart rightly. <code>6.2</code> does have some issues with email processes dying.
Q: I am very interested in the trace log section. Are you going to send the PPT from today?	Unfortunately, I don't have much info on the slides about trace logging. I have spent about 2 hours doing just logging and reading logs in previous talks and would be willing to do that for Vivit.
Q: any tool which can monitor DB, Web and Application performance, recommended? Q: Any specific monitoring tool for HPSM?	1. OVO works quite well 2. Have had good luck with a combination of OVO and site-scope
Q: does purging and archive add to the performance?	Yes, typically it will. Most customers are happy with 2 years of historical data as long as the older data is available for reporting.
Q: Tried to run with 100 threads on a 64 bit. As soon as we had issues with performance and we raised a case with HP ... HP recommend to use 50 only :-) Q: To avoid null table.... As you stated: go to system navigator and fill in the SWL filed mappings manually instead of letting SM do that for you. Worked every time for me now ... since 7.11 at least.	1. I would also recommend 50 threads per process as a max. You could get away with more for your ESS users if you route them to a specific servlet.

<p>Q: is applying windows OS patches on SM server advisable A: --unanswered-- Q: Any way of monitoring background processes ? A: --unanswered-- Q: Thanks Eric & everyone attending A: --unanswered-- Q: hi eric what is significance of tables like- SMSQL1E00446D2228FD & why Sm is not able clen them up A: --unanswered-- Q: How Disaster recovery is created for HPSM any best practices for the creating it A: --unanswered-- Q: hi eric what is significance of tables like- SMSQL1E00446D2228FD & why SM is not able clen them up A: --unanswered-- Q: Thanks A: --unanswered-- Q: hi Eric I have a query from topic databse tuning how to avoid newly added fileds mapped to NullTable N1 A: --unanswered--</p>	<p>1. Yes 2. Best to run anubis process to restart them automatically. It will restart anything you start in the system.start command 3. SMSQL1E...question: These are left-over tables from failed SQL to SQL or dbdict column remappings. They can typically be deleted. 4. Disaster recovery--too long a topic to discuss here 5. Null table issue: Best practice is to add the columns to the dbdict, then edit each column individually and assign sql names and tables to them, don't let the mapping engine do it if you are running into n1 table issues.</p>
<p>Q: What is the difference between unsued and free, if it's unused doesn't that mean it's free? A: --unanswered-- Q: You mention anubis, we run into a problem where linker appears to be running but is not actually doing anything. Would anubis pick up that situation and kill the linker process and restart it? A: --unanswered--</p>	<p>1. Shared memory Unused Space is memory that has never been allocated. Free space is memory that has never been allocated PLUS memory that has been freed to be re-used. 2. Anubis would probably NOT pick up that linker has stopped and re-start it.</p>
<p>Q: can you show us that URL again please?</p>	<p>http://www.t1shopper.com/tools/calculate/</p>
<p>Q: We recently found out that if you're in Horiz LB environment, you must have all your LB servlets offline when adding a field and then restart afterward to avoid the NULL Table issue. Q: Wanted to point out that it's also important to make sure you provide enough JVM memory per Tomcat JVM on the Web Tier side...default of Min32M/Max64M is not sufficient in large environments where most users are comign in through Web Tier.</p>	<p>1. I have not heard of that before and have not seen that behavior, but anything is possible! 2. Excellent point. I'm going to include a section for Tomcat memory in my presentation. I have had good luck with an initial pool of 256 and max pool of 1024</p>
<p>Q: Does horizontal scaling support the fail over of secondary server only ?</p>	<p>1. I'm a little unclear on this question, but horizontal scaling includes multiple machines by definition. It does not include any fail-over for the load balancer portion of SM, I would recommend that be done on two really small machines failed over by MS Clustering or F5 or Veritas clustering. Horizontal scaling can have a large amount of app servers, so true 'fail-over' is not done like what would happen in a vertical scaling mode where you would fail the entire machine over to a new server.</p>
<p>Q: Running Reports affect the application performance , what is your recommendation for reporting Q: How is the failover achieved in the Horizontal scaled example you displayed</p>	<p>1. Recommend having different reporting server, or at the very least, a different reporting db so tables are not accidently being locked up by reporting queries. 2. Failover is not achieve for the load balancer portion, but if a machine goes down, the other two machines are still up to receive requests. I will include a full fail-over solution in my updated presentation.</p>
<p>Q: some tables like msglog, syslog... have a considerable increase ¿they must be purged? ¿how can help those tables to find some unexpected behavior?</p>	<p>1. These tables may be purged, but they are not being read from, so there are no queries being done on those tables. There's no benefit to purging them except to free space.</p>
<p>Q: do you have to use the IR Expert? Q: Everything appears to be related to a windows setup, do you have anything on Unix?</p>	<p>1. IR expert does not have to be used. Either delete the keys from the dbdicts or use the ir_disable:1 parameter. If you are certain you will not be using it, remove the keys. Ir_disable:1 is meant for upgrades and I'm not sure how it will be supported in the future. If unsure, delete the keys. 2. Windows and UNIX are similar for hardware requirements and configuration.</p>