ESRP Performance

Silverton Consulting, Inc.
StorInt™ Dispatch

This dispatch restarts SCI's series on Microsoft Exchange Solution Review Program (ESRP)[1] performance results and reports on the over 5000 mailboxes results category. Prior reports discussed the over 1000 to 5000 mailboxes and the under 1000 mailboxes result categories[2]. As such, to better compare ESRP/Jetstress results SCI reports on both normalized and un-normalized results. For normalized results in this highest-tier category we use operations per 5000 mailbox (5Kmbx). Un-normalized results are in the appendix.

ESRP was never intended to compare subsystem performance but rather as a proof of concept for Microsoft and storage vendors to depict a configuration supporting a given workload. Hence, any comparisons necessarily come with some caveats and may not be real. Nonetheless, SCI feels comparisons can well serve both the vendor and end-user storage community and thus, worth noting.

Latest ESRP V2.0 results

We have added a new ranking for this analysis, which depicts the average database backup throughput across all storage groups. This value correlates moderately to aggregate database transfers per second. (See figure 1).

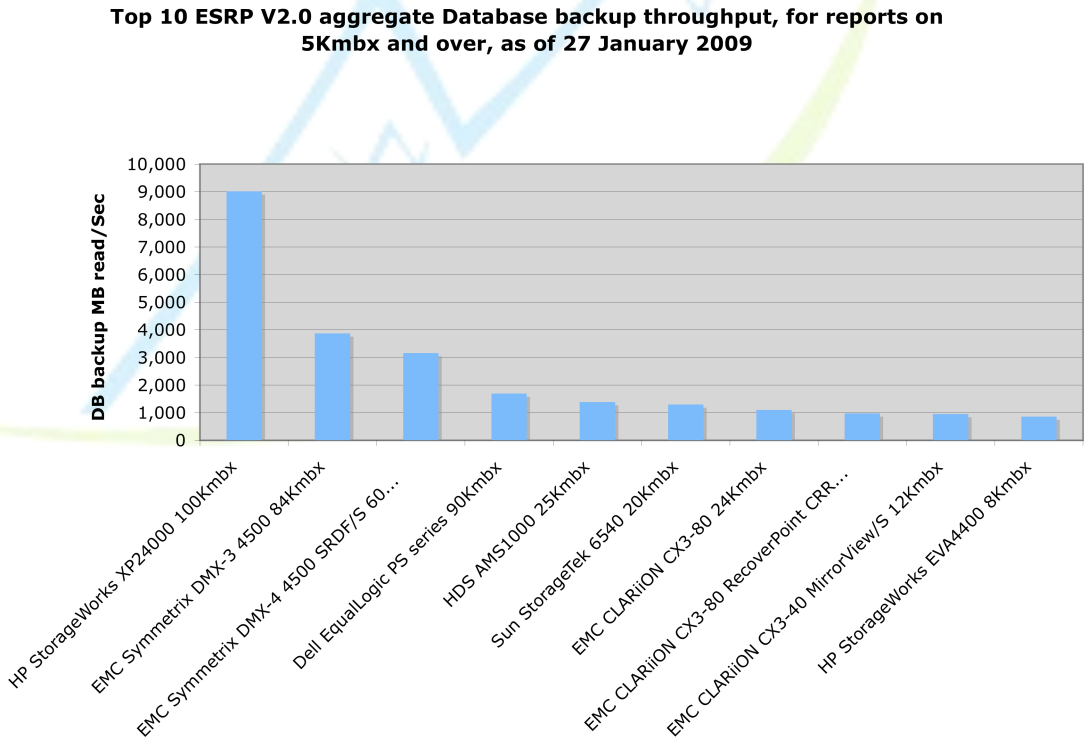

Figure 1 Top 10 database backup results in MB read/second

[1] ESRP results from http://technet.microsoft.com/en-us/exchange/bb412164.aspx, as of 27 January 2009

[2] All of prior SCI ESRP Dispatches can be found at ofhttp://www.silvertonconsulting.com/page2/page2d/storage_int_dispatch.html

As can be seen from the above chart, the new report on HP XP24000 dominates this category. Partly HP's commanding result is due to the overall number of mailboxes being serviced. However, Dell's Equal Logic iSCSI result for 90K mailboxes only came in at number four. Also the number ten result, for HP EVA4400 supports only 8K mailboxes and stands almost as well as the other results supporting many more mailboxes. A couple of caveats worth noting here database backup performance can be impacted by

- Number of disk drives in a configuration
- How message store databases are split across those spindles
- Subsystem RAID level

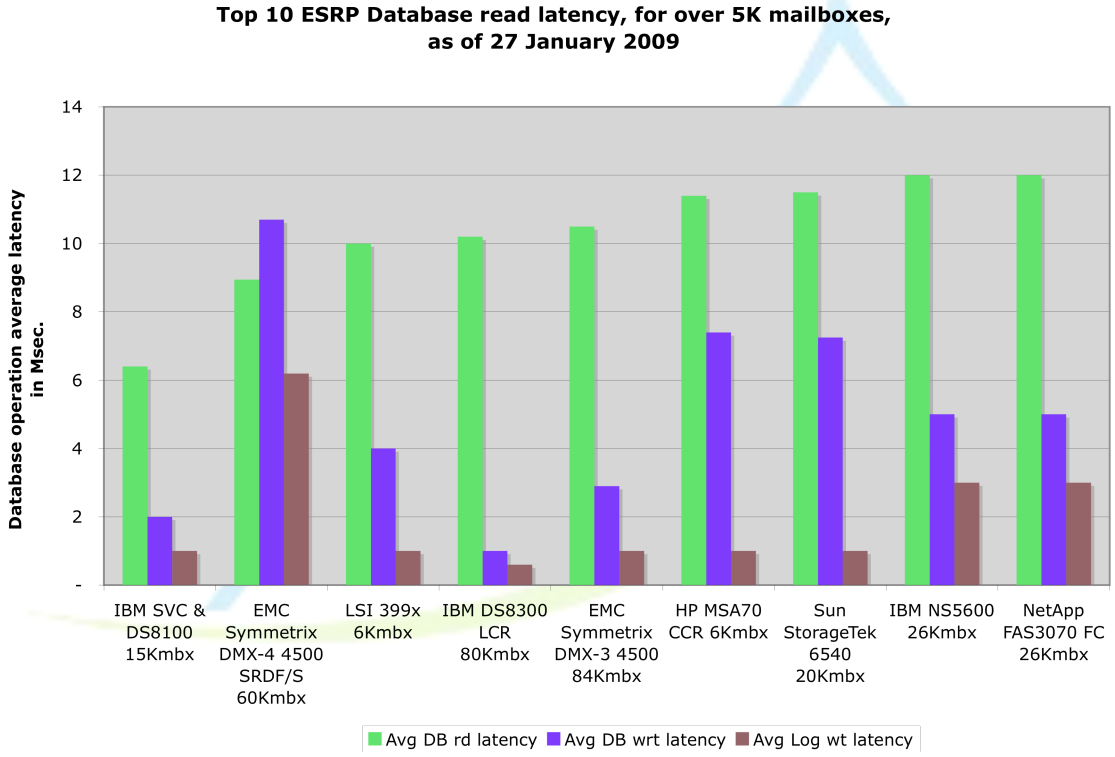

Figure 2 Top 10 ESRP results database operations by latency

We add another new chart for this analysis showing the ESRP reported database latencies for read, write and log write operations. SCI and others feel the read latency metric best shows what an end-user experience would be from a subsystem configuration. The list is sorted by read latency. A couple of considerations to note:

- While read latency is unaffected by replication mode, write and log write latency can be seriously impacted by how the Exchange database is replicated. For example if one examines the EMC SRDF/S in the number two position, its write latency is pretty high. However if one considers that SRDF/S was active this means the data has to be written to the secondary subsystem in parallel to being written to the primary subsystem and as such its write latency does not look that bad.

- There are a couple of ways to impact or game this value. One easy way is to reduce the overall load on the storage. As ESRP reports are intended to show a viable performing solution to handle a simulated user workload we assume that these products are all optimizing cost and performance, so believe this is not an issue here.
- For an ESRP benchmark to be accepted, read latency must be under 20 msecs. Some vendors may try to push read latency out closer to 20msecs in order to support more mailboxes with less hardware. As such, those vendors may not show up well on this top 10 chart.

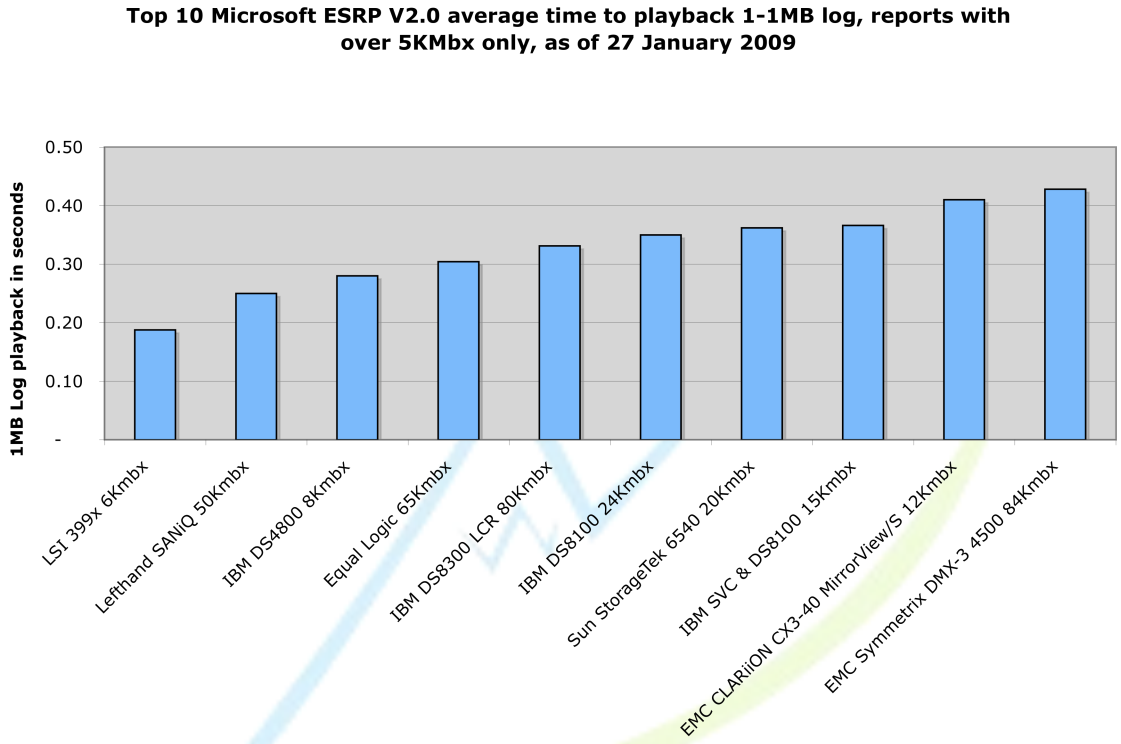

Figure 3 Top 10 ESRP reported log playbacks

Another new chart for this report is the log playback ranking. This is a rather complex workload which encompasses log reading and database reading and updating (or writing). The timings are reported as the average time in seconds it takes to playback or process a 1 MB log. The top 3 systems LSI 399x, Lefthand SANIQ, and IBM DS4800 would all be considered mid range storage subsystems although Lefthand was supporting a heavy workload at 50K mailboxes and was configured accordingly. The range for the top ten subsystems is fairly large over 2X from lowest to highest. It's unclear how one succeeds in this metric other than having fast disk and low latency database operations. Similar to the backup discussions above, some caveats would include:

- Playback performance can significantly be impacted by the number of disk drives.
- How message store databases are split across those spindles also can impact this
- Subsystem RAID level may also impact playback performance
- Replication type may also impact log playback performance

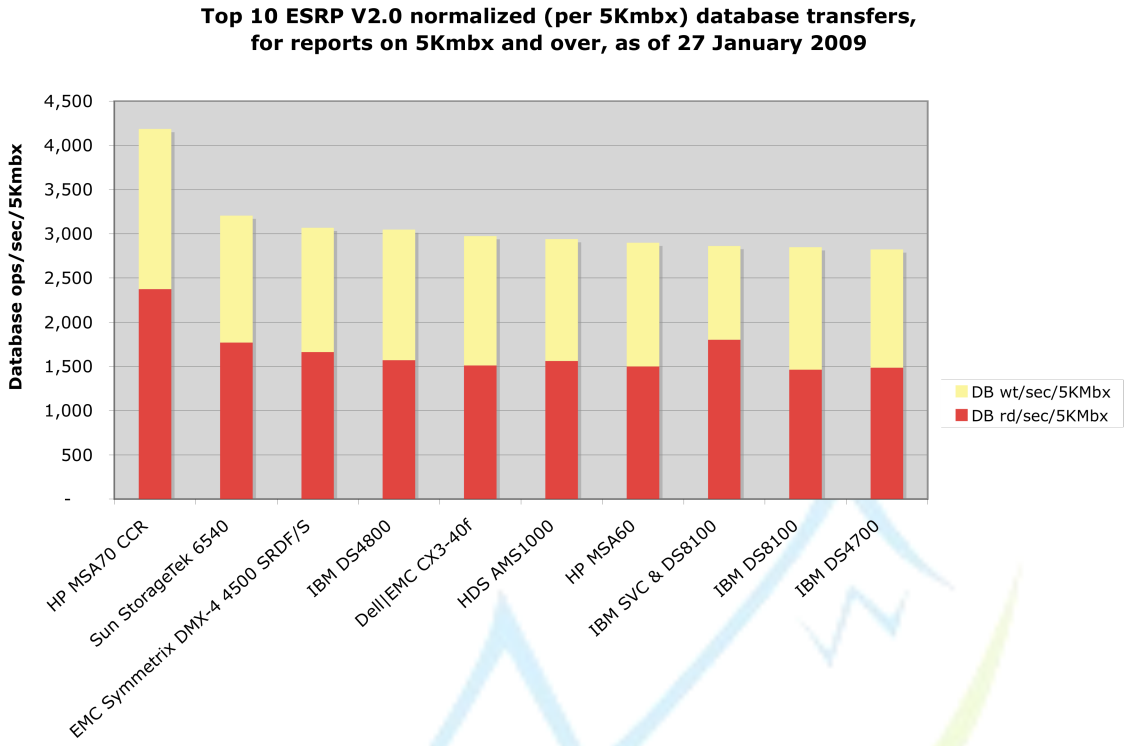

Figure 4 Top 10 normalized database operations/second

Finally, we now turn to overall database transfer results and the Top 3 normalized ESRP/Jetstress results belong to HP MSA70, Sun StorageTex 6540 and EMC Symmetrix DMX-4 4500. A few considerations are warranted on normalized results:

- Normalized results do not often scale well. Although four of these results were for 20,000 mailboxes or over, (Sun 6540 at 20KMbx, EMC Symmetrix at 60KMbx, HDS AMS1000 at 25KMbx, and IBM DS8100 at 24KMbx) the top result from HP supported only 6000 mailboxes and may not scale much beyond that quantity of mailboxes.
- One surprise here is the close running of everyone behind the top result and may be an artifact of the ESRP benchmark striving to generate equivalent workloads per user mailbox. However, the over 5K mailbox tier is the only category that shows results this close to one another. (see prior ESRP StorInt Dispatchen for more information).

Conclusions

From our perspective, ESRP results in this over 5K mailbox tier are getting more competitive. There were a number of new ESRP results in this category over the last 9 months, and at least 4 over the last quarter. Seeing HP (an OEM version of HDS's top product), EMC Symmetrix and IBM DS8300 running the same performance tests is a good indicator of their willingness to show their products in the best light as well as high customer interest in Exchange solutions. Probably these vendors don't see their individual results as entirely comparable and arguably they may have a point, but we would differ with them on this assessment. Moreover, seeing SAS and iSCSI results compete in this top-tier at least on various metrics indicates these interfaces can provide some competition to FC storage in mission critical applications.

ESRP/Jetstress results are inherently difficult to compare. Nonetheless we believe Exchange results provide a unique real world benchmark and deserve some comparison so that the public can make properly informed storage purchases. Our next ESRP/Jetstress report will return to the 1K to 5K mailbox tier. We continue to welcome any feedback on how to do better.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community

Silverton Consulting
Strategy, Storage & Systems
P.O. Box 492|Broomfield, CO 80038-0492|
720-221-7270|www.silvertonconsulting.com

Appendix

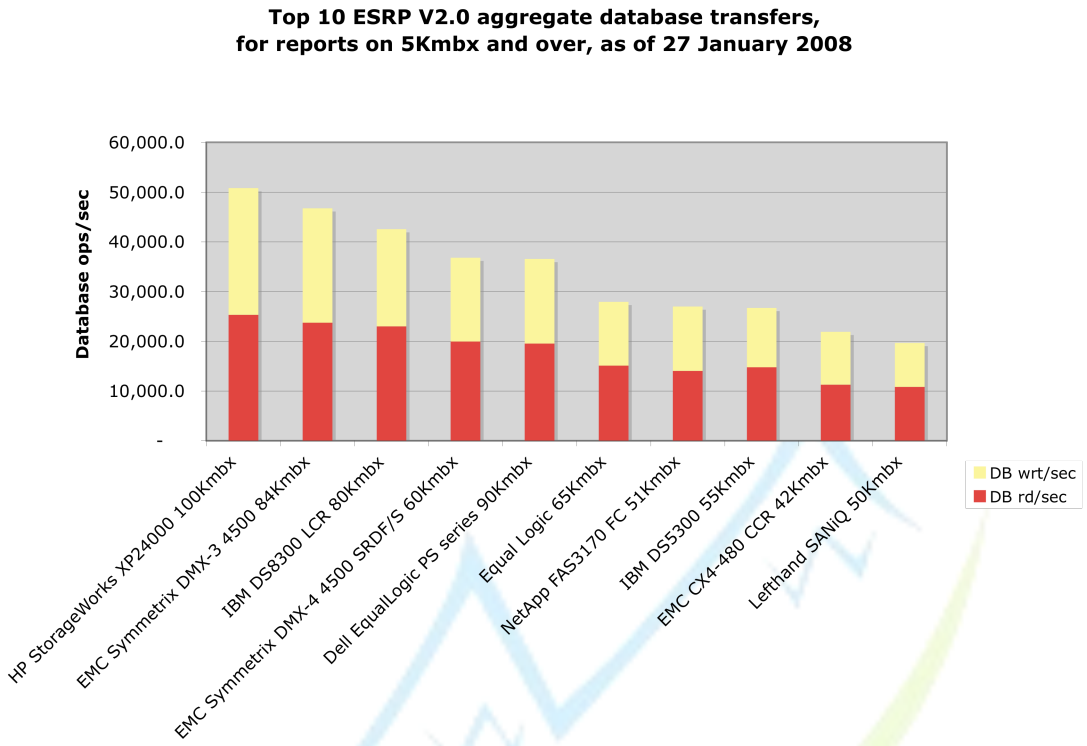

Figure 5 Top 10 Aggregate database transfers per second

Aggregate or un-normalized database transfer results are highly correlated to number of mailboxes in service and as such, are relatively less useful metrics in our opinion. Nonetheless, it's probably no surprise that the top results belong to HP, EMC and IBM. Aside from the iSCSI results, this ranking is similar to one ranked purely on number of mailboxes.

Silverton Consulting
Strategy, Storage & Systems
P.O. Box 492|Broomfield, CO 80038-0492|
720-221-7270|www.silvertonconsulting.com