

We now turn to analyzing Exchange Solution Review Program (ESRP)¹ results in the 1001 to 5000 mailbox category. Our January report² discussed the over 5000 mailboxes category, and our next ESRP report will discuss the under 1001 mailboxes result category. To better compare ESRP/Jetstress results SCI reports on both normalized and un-normalized results. For normalized results in this mid-tier category we use operations per 1000 mailbox (1Kmbx). 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.

Latest ESRP V2.0 results

We now use a ranking for this analysis depicting the average database backup throughput across all storage groups, which roughly corresponds to an IT or datacenter view of Exchange performance. This value correlates moderately to aggregate database transfers per second. (See figure 1).

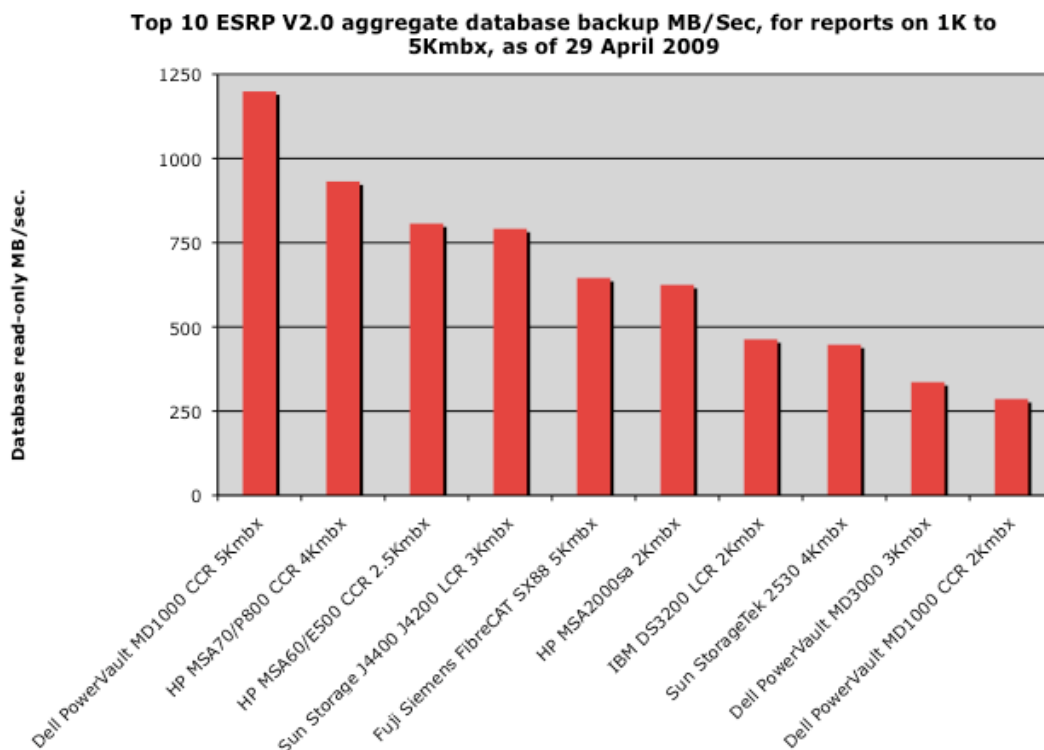


Figure 1 Top 10 database backup results in MB read/second

¹ ESRP results from <http://technet.microsoft.com/en-us/exchange/bb412164.aspx>, as of 29 January 2009

² Available at <http://www.silvertonconsulting.com>

As can be seen from the above chart, the Dell's MD1000 report wins this category. Partly Dell's result is due to the overall number of mailboxes being serviced. However, Fuji Siemens result, also for 5K mailboxes only came in at number five. 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

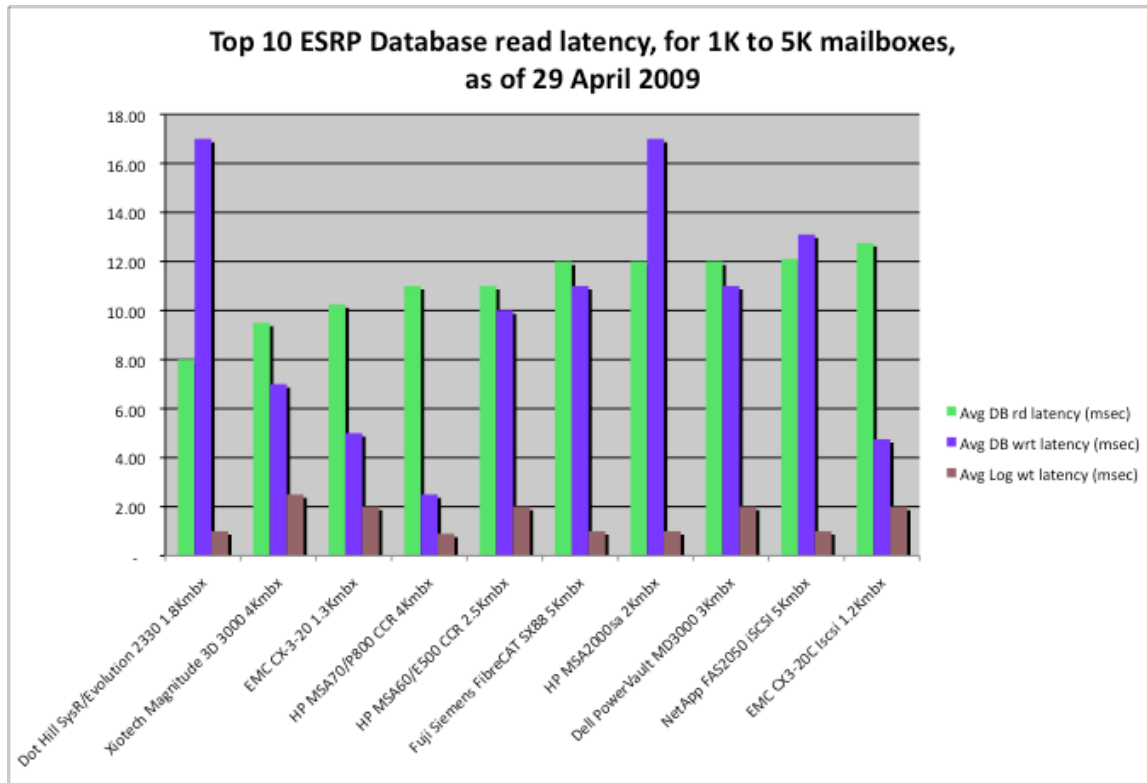


Figure 2 Top 10 ESRP results database operations by latency

This chart shows the ESRP reported database latencies for read, write and log write operations. SCI and others feel the read latency metric best shows what an Exchange end-user experience would be from a subsystem configuration. The list is sorted by read latency. Not sure why Dot Hill's write latency is so poor here but it could be an artifact of their caching activity. A couple of other 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.
- 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

The log playback is a rather complex workload that 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 Dot Hill SystemR, and the two Dell Power vault MD1000 results would all be considered small to medium range storage subsystems. The range for the top ten subsystems is fairly large over 3X 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:

- The number of disk drives can significantly impact playback performance.
- 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

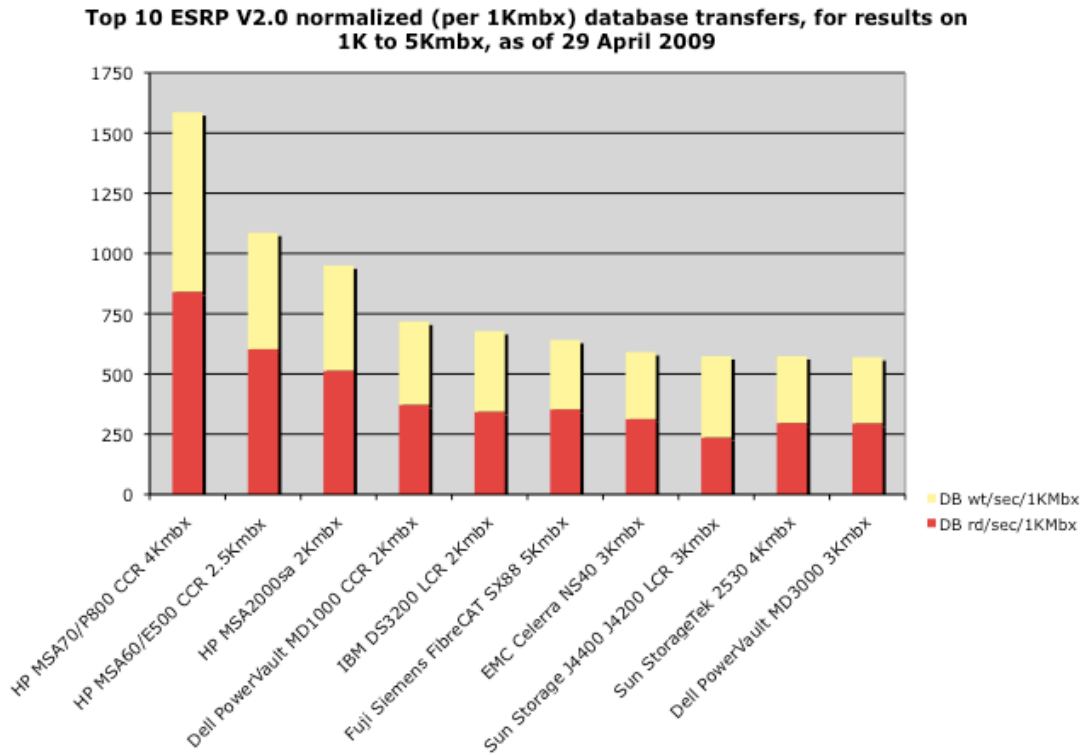


Figure 4 Top 10 normalized database operations/second

Finally, we now turn to overall database transfer results and the Top 3 normalized ESRP/Jetstress results belong to HP MSA70, HP MSA60, and HP MSA2000sa. One consideration is warranted on normalized results, namely that normalized results do not always scale well. Although only one of these results were for 5,000 mailboxes (Fuji Siemens FibreCat), the top result from HP supported only 4000 mailboxes and may not scale much beyond that quantity of mailboxes.

Conclusions

From our perspective, ESRP results in this 1K to 5K mailbox tier are heating up since we reported on this category last July. There were a number of new ESRP results in this category over the last 9 months, and at least 2 over the last quarter. For some reason HP seems to have a lock on this category performance. Also the category supports SAS, FC and iSCSI solutions very evenly.

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 and under mailbox tier. We continue to welcome any feedback on how to do this analysis 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

Appendix

Top 10 ESRP V2.0 aggregate database transfers, for results on 1K to 5Kmbx, as of 29 April 2009




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. How HP MSA70, a SAS result, dominates this category is difficult to discern.