

This dispatch covers Microsoft Exchange Solution Review Program (ESRP)¹ V3.0 for Exchange 2010. There have been a large number of submissions this past quarter with at least ten new ones in the over 5K mailbox category discussed here. Future reports will cover the 1001 to 5K and the 1K and under mailbox categories. Previous ESRP V2 and ESRP V3.0 analysis reports are available on our website².

Latest ESRP V3.0 results

We start our analysis with Exchange database access latency results. Recall that this chart lists the top 10 database-**read** latencies reported by ESRP for the over 5K mailbox category.

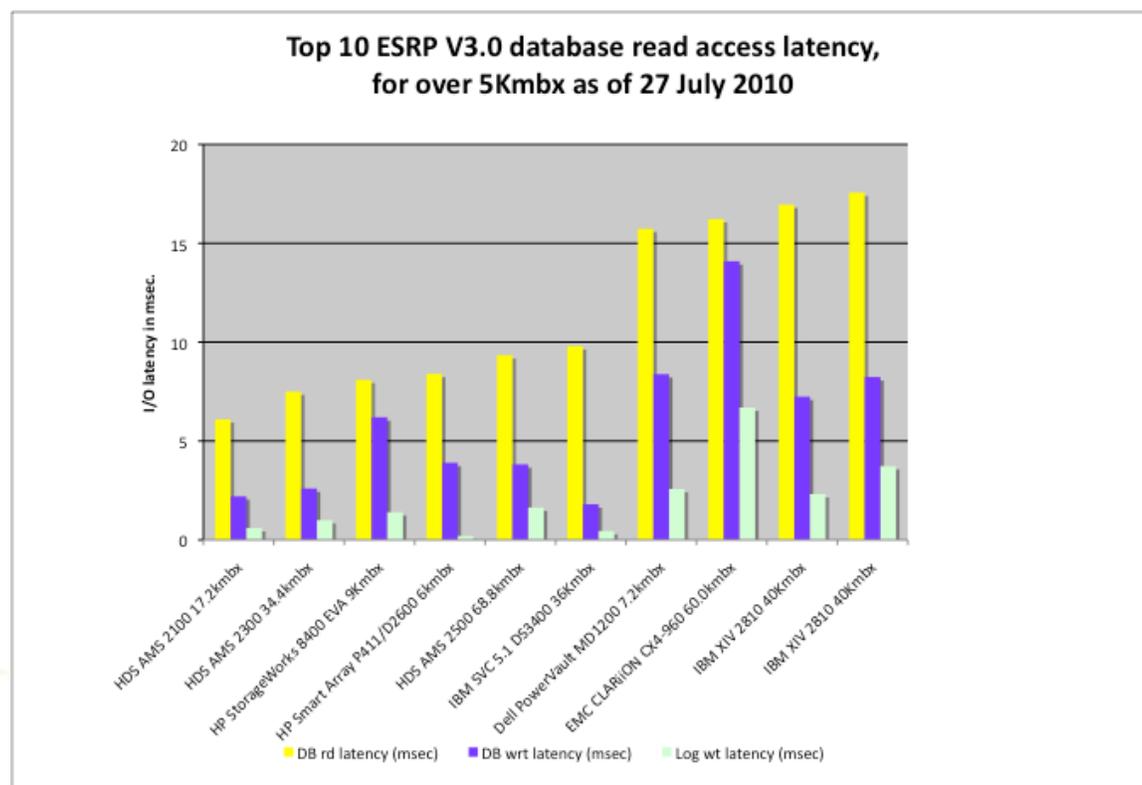


Figure 1 Top 10 ESRP database-read latencies

For database reads the HDS AMS 2100 and 2300 take the top two spots with HP's EVA coming in a close third. However, I find the HP Smart Array results (#4) very interesting.

HP's Smart Array consists of just a bunch of SAS interfaced disks, connected to Exchange Servers. Its log-write latency is almost immeasurable (~0.1msec). As

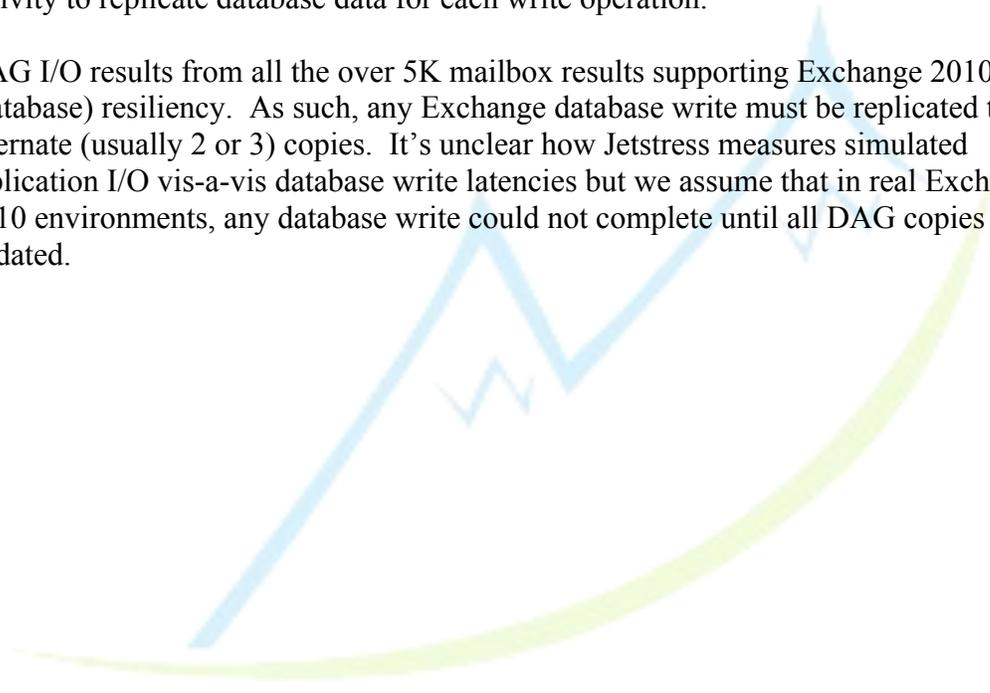
¹ ESRP results from <http://technet.microsoft.com/en-us/exchange/ff182054.aspx>, as of 27 July 2010

² All prior SCI ESRP Dispatches can be found at <http://silvertonconsulting.com/cms1/dispatches/>

log writes are primarily sequential writes and we would expect a JBOD to do well here but this seems too good. That this (sub-)system had excellent database read and log-write latencies implies that both disks and controllers were well tuned for random and sequential I/O - in my experience, a hard thing to do without cache. HP's Smart Array used a mailbox database size = disk drive size, which may have resulted in good access times, but it's unclear why. As a counter example, Dell's submission (#7) also used direct connected SAS drives but had a database size ~2X the size of their disks.

The other odd result in Figure 1 is the variability in Exchange 2010 database write access times. One would think that caching subsystems would accommodate most database writes at high speed. But, we believe this write variability is an affect of Exchange 2010's larger database block sizes forcing more destage activity and simulated DAG I/O activity to replicate database data for each write operation.

DAG I/O results from all the over 5K mailbox results supporting Exchange 2010 mailbox (database) resiliency. As such, any Exchange database write must be replicated to alternate (usually 2 or 3) copies. It's unclear how Jetstress measures simulated replication I/O vis-a-vis database write latencies but we assume that in real Exchange 2010 environments, any database write could not complete until all DAG copies were updated.



Next we turn to database transfer counts. As ESRP mailbox counts for the over 5K mailbox category span such a wide spectrum we have elected to normalize database transfer counts to accesses per 1,000 mailboxes (1Kmbx).

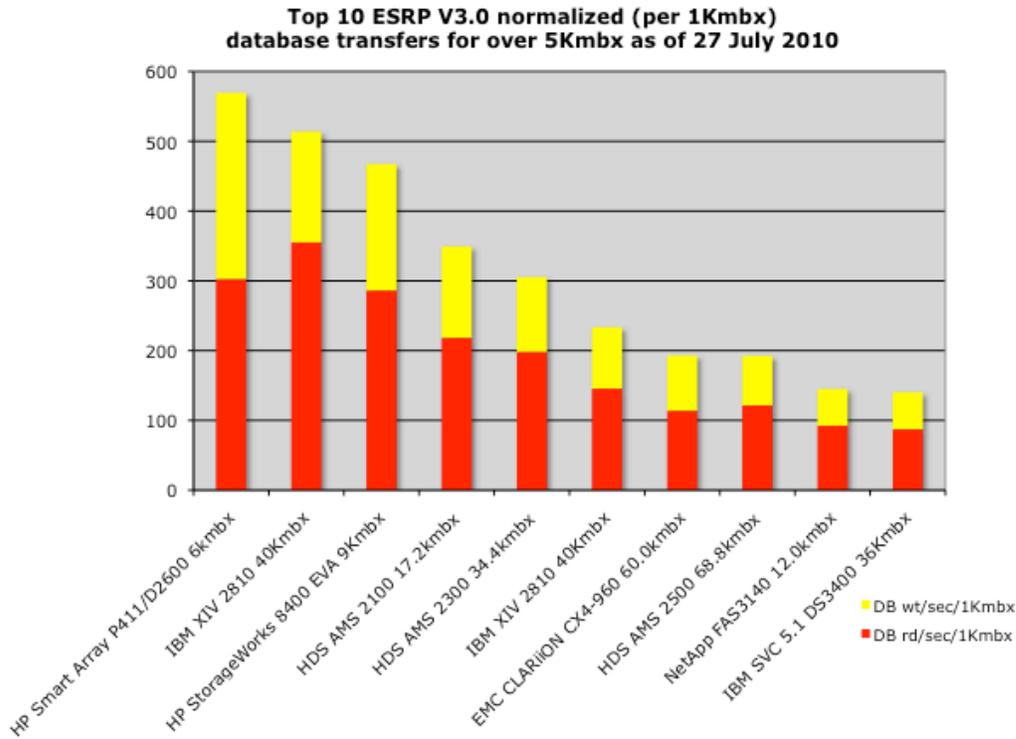


Figure 2 Top 10 normalized DB transfers per second

The #1 result was the same HP Smart Array discussed above an HP Proliant server SAS connected to an D2600 disk array. In addition, there were two IBM XIV submissions (#2 and #6) at the same mailbox count (40K) with substantive differences between the two being the mailbox size (3- vs. 1-GB for slower), drive capacity (1- vs. 2-TB for slower), and storage used (40% vs. 88% for slower). As discussed in prior reports, the read database transfer counts are significantly higher than the write transfers, in some cases almost 2X the rate.

A couple of caveats for normalized results:

- Normalized results may or may not scale much beyond the reported mailbox counts. For example, the #1 result supported 6K mailboxes and may not support much more than that.
- Normalized results can be impacted by over provisioning. For example, the #2 result only used 40% of its storage for Exchange services allowing it to use more spindles than necessary for the workload.

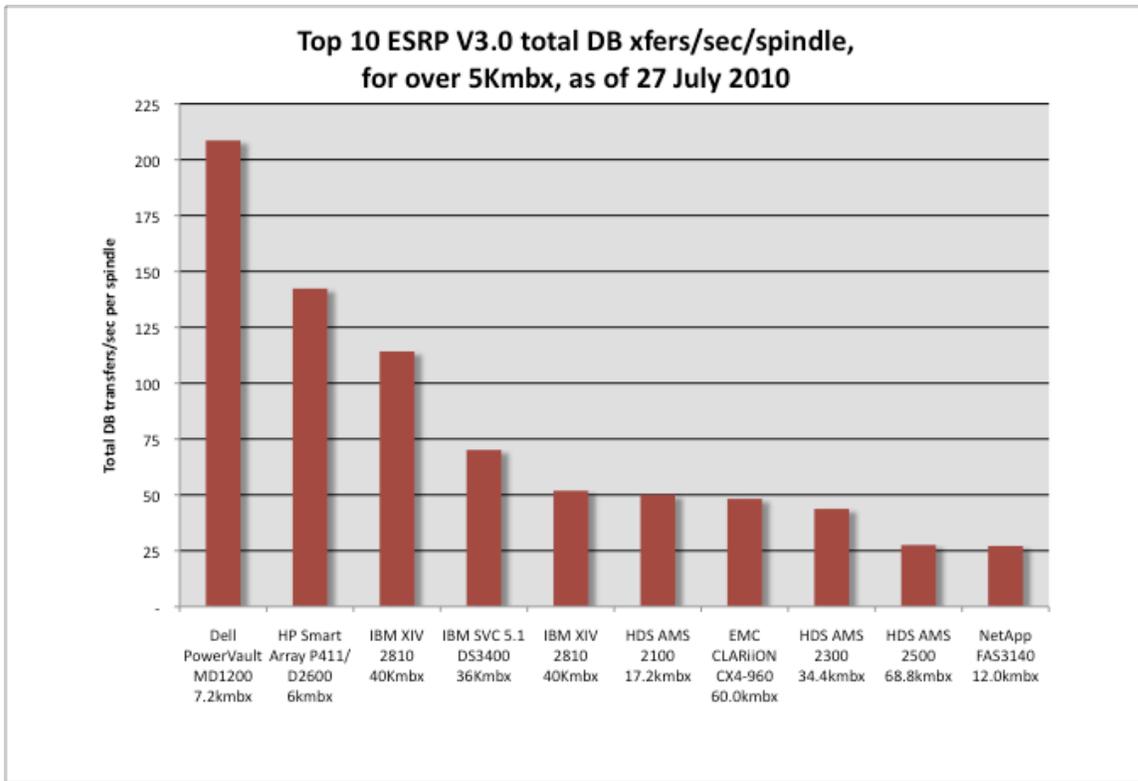


Figure 3 Top 10 data base transfers per spindle

Speaking of over provisioning, another way to look at Exchange storage performance is to normalize it over the number of disk drives or spindles used for the configuration. Figure 3 above shows the total number of (read and write) database operations per second per drive done by each subsystem. Here one can see the same two IBM XIV submissions in the same order discussed for Figure 2.

Some caveats for database transfers per spindle results:

- Drive speed can help one do well on this metric, i.e. 15Krpm drives can perform better than 7.2Krpm drives and four of the top six performers used 15Krpm drives (both XIV results used 7.2Krpm drives).
- Drive over provisioning usually reduces one’s performance on this metric. However this was not evident with the XIV placements (the over provisioned XIV did better than the normally provisioned one).

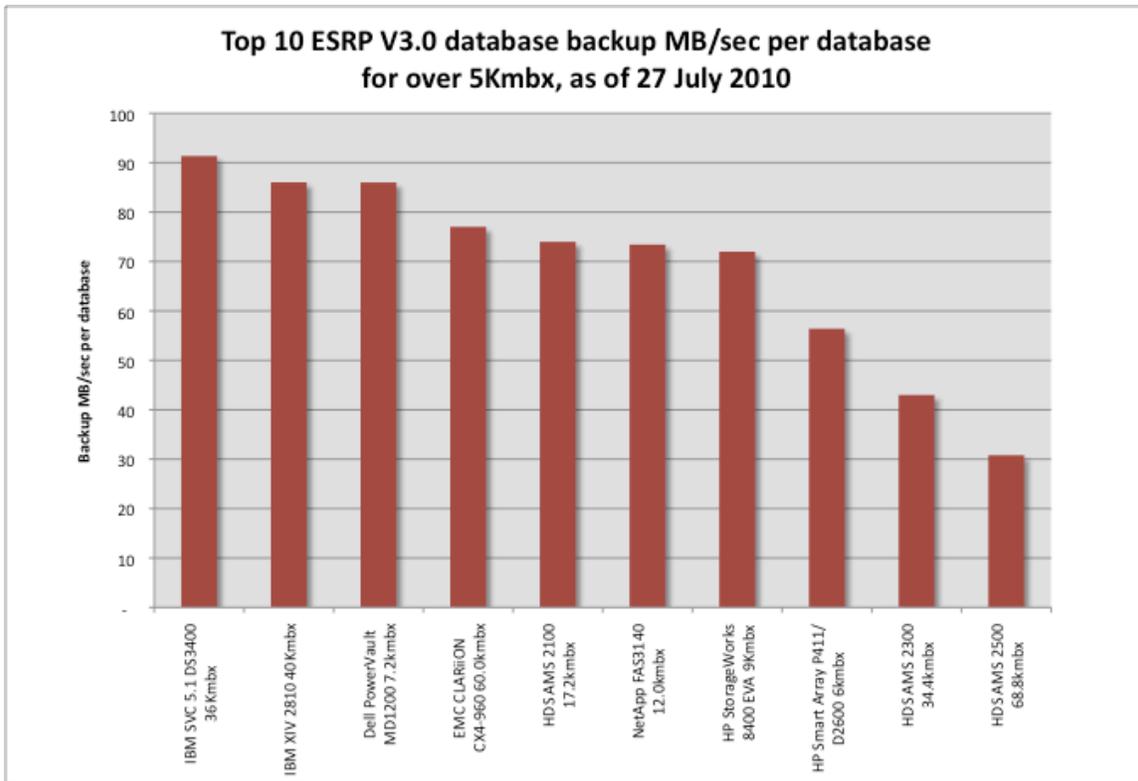


Figure 4 Top 10 database backup activity in MB/second/database

Next, we show the top 10 Exchange backup throughput rates in MB/sec/database. The #1 and #2 positions went to IBM with SVC and XIV. Tied for #2 spot was Dell PowerVault, the other SAS connected disk system.

With Exchange 2010 and mailbox resiliency through DAGs, database backup activity no longer seems as important. In fact, there was at least one submission that didn't even report on this metric. However, there are many valid reasons for database backups and we continue to believe that there will be an ongoing need for mailbox backups. As such, reporting on backup speed need to be re-instated and preserved.

Conclusions

ESRP results for the over 5K mailbox category are always difficult to analyze due to the wide span of mailbox counts (from 6K to almost ~69K for current submissions). That said, with the limited submissions to date, it appears for smaller mailbox counts, a properly configured SAS direct connected storage system may perform well enough. Above that is subject to some debate but more results should help clarify this.

Nonetheless, it's still early in ESRP v3 history. To date there have been only 12 submissions in this category and just over 20 overall (with reports available). Even so, we were surprised to see this many, since our last report only showed 4 results for all categories.

Finally, ESRP/Jetstress results seem designed to be difficult to compare but merit the effort. Thus, we strive to improve our analysis with each report. As always, feel free to contact us with any ideas on how to improve. In that regard, our contact information can be found below or on our website at SilvertonConsulting.com.

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