TSMPoint

USER SOFTWARE WORKS

Executive Summary

After a successful SharePoint® implementation, if the users realize its potential (versioned storage, etc.), the volume of stored documents could rise exponentially. This is the so called „SharePoint boom". Documents and their indices are stored as BLOBs in an MS SQL database. The database rapidly inflates, which could impact the overall performance of the system and can be also challenging for the database administrators (backup and database maintenance)

Using TSMPoint, only document metadata is stored into the database, documents can be directed into Spectrum Protect™, eliminating database inflation. Documents get directed to the database or to Spectrum Protect based on their size. Small documents get stored into the database, documents bigger than a configurable size limit get directed to Spectrum Protect.

Benefits

- The decreased size of the SharePoint database allows a higher percentage of it to be kept in memory, increasing its performance.
- The smaller SharePoint database size also speeds up its backup and restore.
- Existing SharePoint databases can be migrated using built-in SharePoint commands.
- Documents stored in Spectrum Protect can be moved back to the MS SQL database with standard methods, should the need for system restructuring arise.

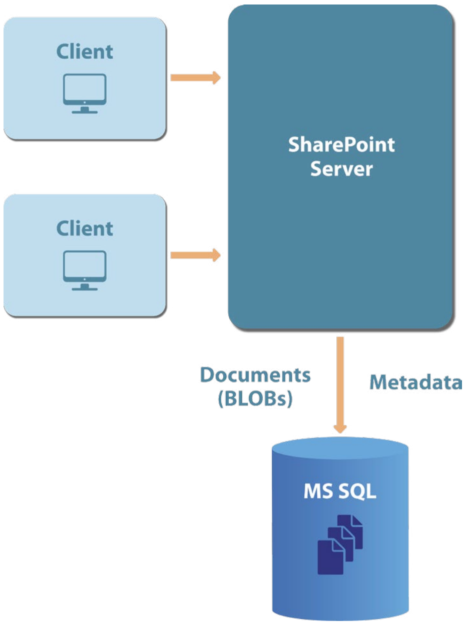

Figure 1. Without TSMPoint

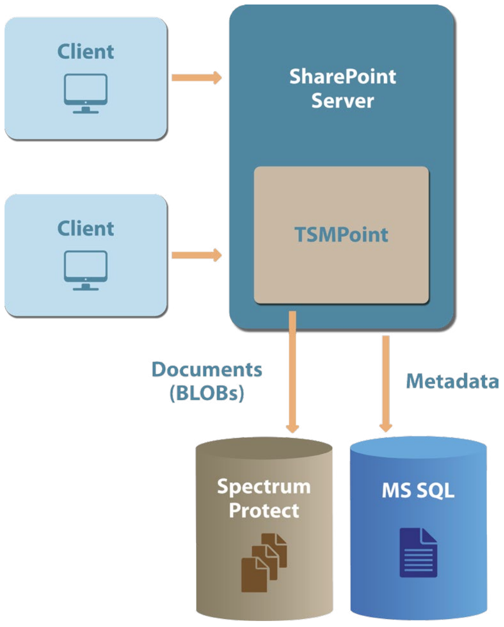

Figure 2. With TSMPoint

SharePoint is a registered trademark of Microsoft Corporation in the United States and/or other countries. Spectrum Protect is a trademark of International Business Machines Corporation registered in many jurisdictions worldwide.

Test Results

Performance and validation tests were run in IBM Innovation Center Budapest on two IBM System x3750 m4 (8752) servers connected to an IBM FlashSystem v840 storage. One of the x3750 servers run the virtual servers (AD, SharePoint, MS SQL), IBM Spectrum Protect (formerly TSM) has been deployed on the other server over a Linux OS. The servers were interconnected via a 10Gb ethernet network, the storage was connected via a 8Gb fibre channel network. To get comparable results (in terms of read and write performace), both Spectrum Protect and MS SQL used a similar LUN on the same storage.

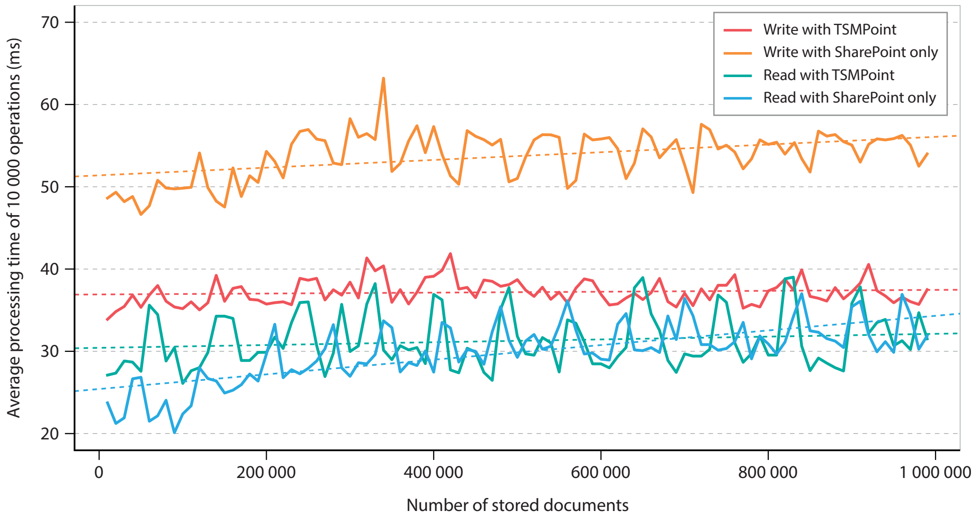

Our experience in the course of performance tests was that, when using TSMPoint, the average storage time of random-sized documents, between 20 kBytes and 1 Mbytes, was significantly less than for SharePoint MS SQL database in a native SharePoint solution in the form of BLOB.
Two additional conclusions can also been drawn on the basis of the trend curves. On the one hand, in the case of a pure Sharepoint solution, the storage speed apparently increases as a function of the number of stored documents, while this growth is insignificant when TSMPoint is used. One the other hand, although initially a small database size takes less read time in the native SharePoint solution than realized by means of TSMPoint, this order varies according to the growth of the number of stored documents, and the TSMPoint solution is the quicker one after 700.000 documents have been stored.

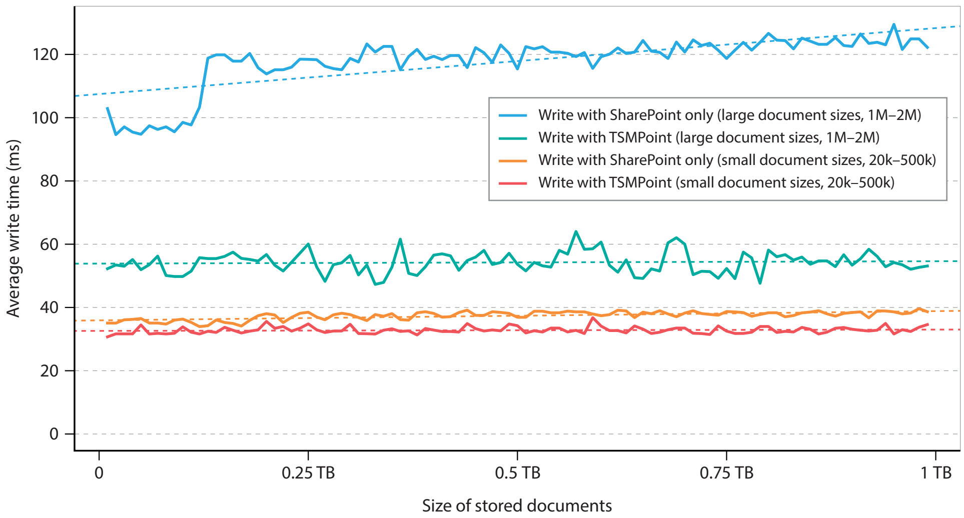

The second figure shows that document storage time is significantly influenced by their size. In the native SharePoint environment, the average storage time of random-sized documents between 20 and 500 kByte (average size of 270 Kbyte) was about 35 ms, with larger documents, those between 1 and 2 Mbytes (average size, 1.5 Mbytes), the average storage time was about 3.5 times greater than the previous value. The increase in the storage time is only about 1.5 times greater when TSMPoint is applied.

In summary, the application of TSMPoint not only simplifies the operation of the SharePoint environment, but can (also) result in performance gains that are clearly observeable to the enduser in cases where average document size is larger size (>500 kbyte) and/or a large number of documents are being processed, as long as the storage speeds of MS SQL database and TSM are commensurable.

Disclaimer: as mentioned in the test environment description, read and write performance is comparable only if the storage performance is similar in the IBM Spectrum Protect and MS SQL environments.