Ceralon Acumen Performance summary

for single server environment on

on HP BL20p G3 and BL460c

Acumen is an integrated e-Discovery solution that delivers import, review, production and export functionality on a scalable multi-tiered server architecture. The following tests, performed at HP's PTAC, compare HP Blade server performance to HP DL380 results that were captured in earlier testing.

Contents

Acumen Import - results

In Acumen's Import module, files are inventoried and documented with custodian information. All metadata, text and MD5 hash values are extracted. Compound files are exploded, de-duplicated and optionally filtered by search terms.

Ceralon Acumen Import on HP Proliant BL20p G3 and BL460c vs. HP DL 380

Performance in megabytes/ hr with repository configurations

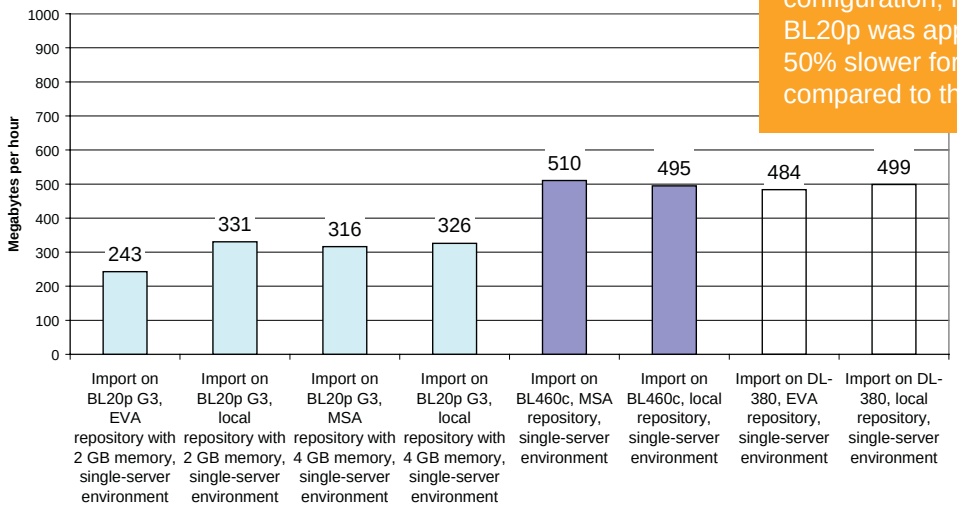

Acumen imported up to 510 MB per hour on the BL460c, using the SAN for repository and database files. A 1 GB Outlook pst file comprised of 23,000+ files took approximately 2 hrs 4 min to import.

The BL460c was approximately 5% faster at Acumen Import than the comparable DL-380 configuration; however, the BL20p was approximately 50% slower for the same test compared to the HP DL380.

Hardware specifications

This sheet summarizes hardware specs of the single-server Blade environments on which Acumen was tested:

HP ProLiant BL20p G3 server

- CPU: 2x 3.06 GHz
- Memory: 2GB, PC2-3200 DDR2
- Local disk: 2x36 HDDs (36GB RAID 1 - local logical drive)
- SAN: EVA6000, two switches (one internal). 2-controllers, 512 cache, 56 drives x 72gb @ 15k (max ~4TB) (two logical drives, each has 200GB with RAID-0). Midway through the trial, switched to MSA 1000 w/512mb cache, 14 drives x 36gb 10k Disks, & Imbedded 2gb fc switch (two logical drives, each has 200GB with RAID-0).
- OS: Windows 2003 SP1, Server edition
- MS Office Professional 2003
- IIS 6.0 with FrontPage extensions
- SQL Server 2000, SP4

HP ProLiant BL460c server

- CPU: 2x Intel Xeon Processor 5160 - Dual core / 3.00 GHz / 1333MHz FSB
- Memory: 4GB of memory, PC2-5300 DDR2
- Local disk: Local 2x146GB SAS HDD drives - (146GB internal logical drive-RAID-1)
- SAN: MSA 1000 w/512mb cache, 14 drives x 36gb 10k Disks, & Imbedded 2gb fc switch (two logical drives, each has 200GB with RAID-0).
- OS: Windows 2003 SP1, Server edition
- MS Office Professional 2003
- IIS 6.0 with FrontPage extensions
- SQL Server 2000, SP4

Detailed specifications:
http://h18004.www1.hp.com/products/servers/proliant-bl/c-class/460c/specifications.html

The HP Proliant BL 460c server performed as well or better in these trials than the HP DL380, while the BL20p lagged behind.

In an attempt to raise performance on the HP Proliant Blade BL20p to the level of the HP DL380, the following hardware changes were made:

- Added memory. Raised memory on the BL20p G3 from 2 GB to 4 GB.
- Changed SAN device from EVA6000 to MSA1000 in an effort to eliminate the second switch required to connect the EVA6000 to the BL20p G3.

Despite these changes, BL20p G3 still lagged behind the DL380 for Acumen Import performance.

Acumen Import - system performance

Data characteristics

Trial data consisted of a 1 GB pst containing 23,797 files, of which approximately half (540 MB) were messages, and the other half consisted of ppt (180 MB), docs (150 MB), pdf (86 MB), xls (30 MB), and other file types.

Capacity utilization

The following charts show server utilization on the HP ProLiant BL460c under load from a 1 GB pst import with Acumen. An import with the file repository and database located on the local disk is compared to an import with the file repository and database located on SAN (MSA 1000).

Memory utilization and average disk queue length during the inventory and compound file explosion phase (first 30 minutes) and again during the de duplication and index phase of import (last 40 minutes). Although total CPU utilization on this 2-dual core CPU system never exceeded 40%, it peaks during the dedup and index phase with the local repository configuration.

Acumen Import on HP ProLiant BL460c
Avg. Disk Queue Length with 1 GB pst

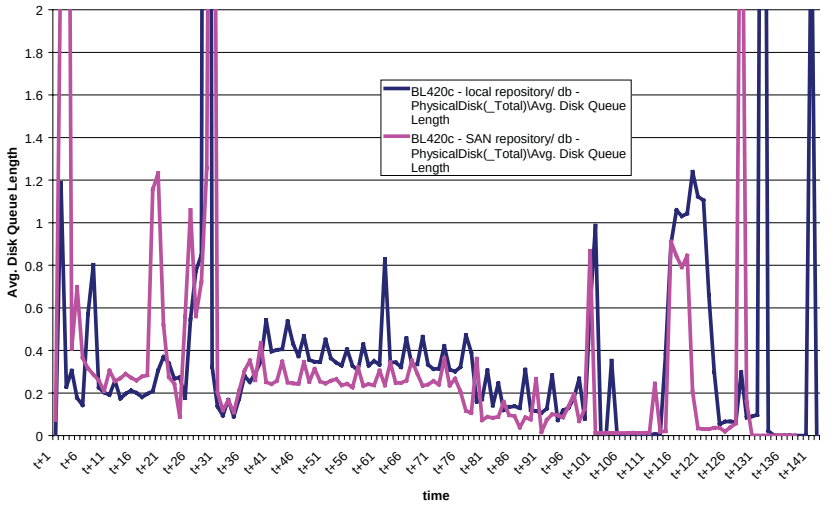

Acumen Import on HP ProLiant BL460c
Memory usage with 1 GB pst

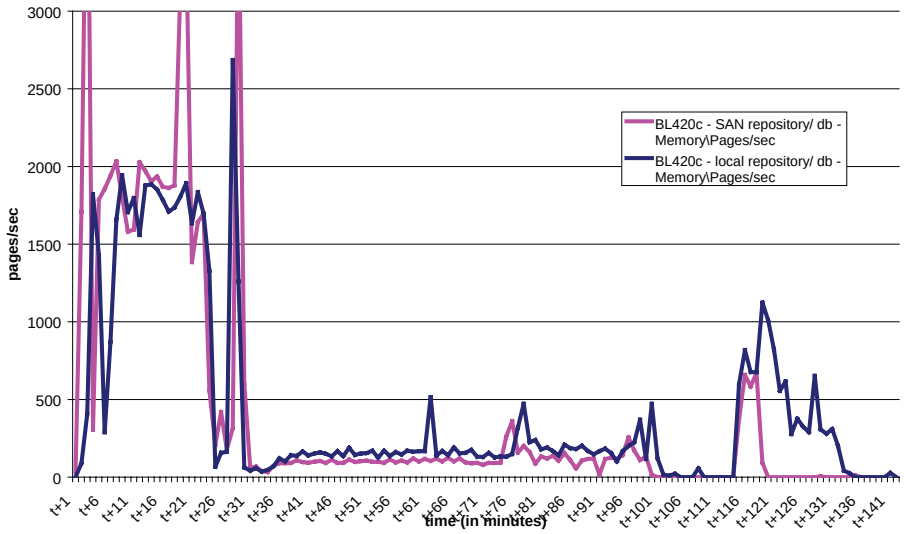

Acumen Import on HP ProLiant BL460c
CPU utilization with 1 GB pst

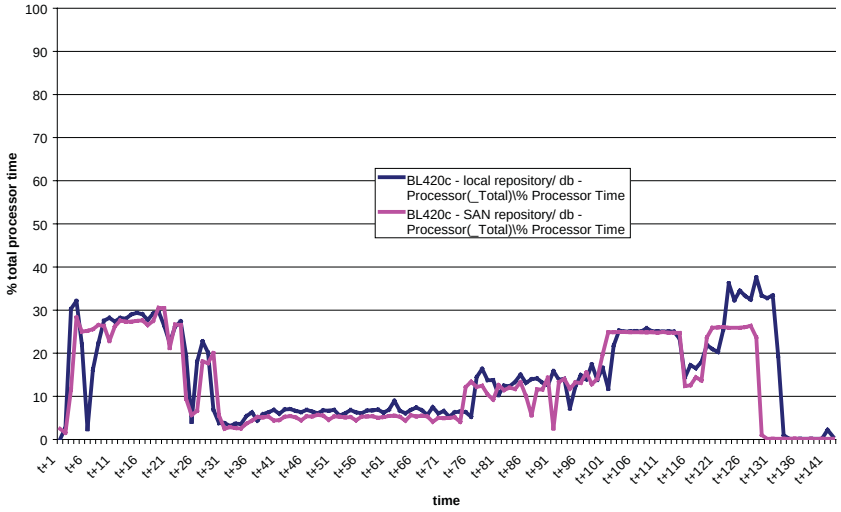

Acumen Production - results

Acumen produces responsive documents for either native file or image file export. The results that follow measure Acumen's performance with the production process running on a single server; however, Acumen may be configured to run production processes on multiple servers simultaneously.

Acumen tiff'ed up to 4400 pages per hour for a set comprised of mixed file types.

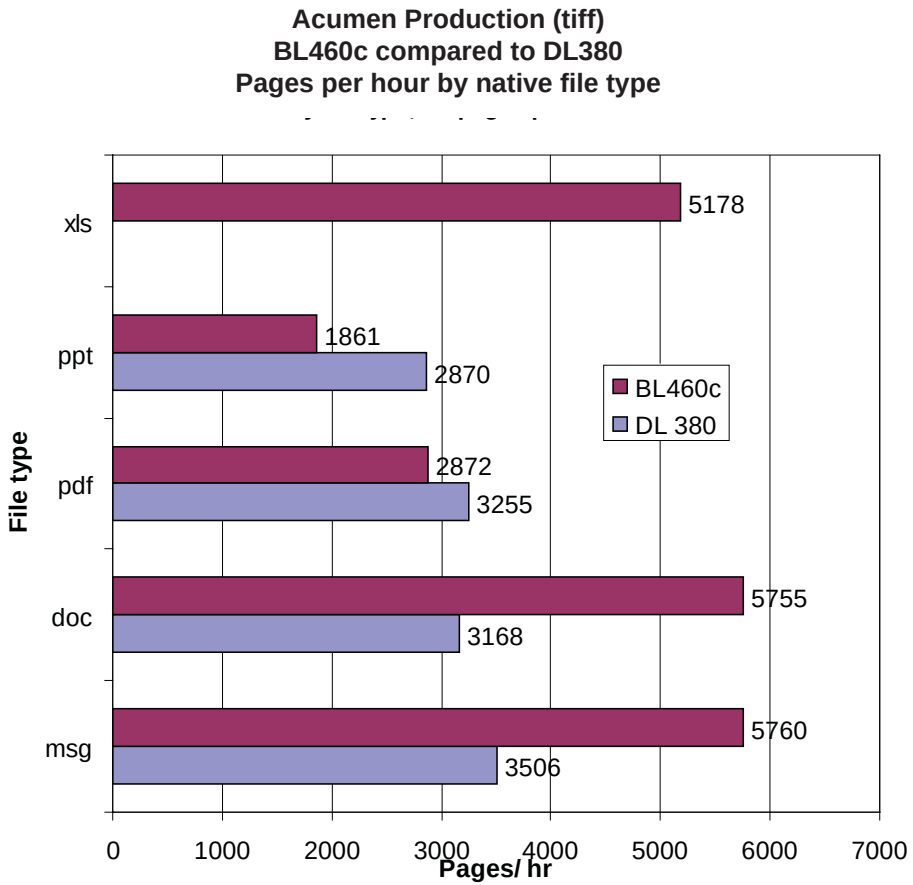

Data characteristics

Files were tiff'ed using Black and White Group 4 compression, although this setting may be changed in Acumen at the time of production to use LZW or other compressions. Production performance varied based on the file types contained in the production set. See graph above for results.