


Pure Storage FlashArray Benchmarking



Perforce Consulting
April 2014

OVERVIEW

This report shows two of the standard benchmarks from Perforce to quantify the performance characteristics of the Pure Storage FlashArray.

To perform this benchmark, Perforce utilized a FA-420 from Pure Storage. The results of the benchmarks are presented within this report.

The Pure Storage FlashArray is configured with multipath and NOOP scheduler. A single 8 Gig dual port fiber card was used to attach the Pure Storage to the server. No FC switches were utilized.

REFERENCES

The following documents were referenced for specific configuration guidelines and recommendations:

- <http://support.purestorage.com/entries/23731087-Linux-Initiator-Settings-Best-Practices>

HARDWARE/SOFTWARE

The benchmarks focus on “read” and “write” performance of the Pure Storage FlashArray as it applies to Perforce applications.

| Binary | Perforce Software |
|---------------|---------------------------------------------------|
| P4 | Rev. P4/LINUX26X86_64/2014.1/807760 (2014/03/18) |
| P4D | Rev. P4D/LINUX26X86_64/2014.1/807760 (2014/03/18) |

The following tables list the hardware configuration of the Pure Storage devices used in this analysis:

| Description | Pure Storage FA-420 Series Memory Array |
|--------------------|------------------------------------------------------------------|
| Model | Pure Storage FlashArray FA-420 with Purity version 3.4.2 |
| Storage Shelves | 2 shelves, each with 22 cMLC 238GB drives and 2 2GB NVRAM drives |
| Controllers | 2 FA-420 controllers interconnected via 2x 56 Gb/s InfiniBand |
| FCP | 4/8Gb/s Fibre Channel x8 |

The following tables list the server hardware configuration used in this analysis:

| Description | P4D Server Hardware Specifications |
|--------------------|-----------------------------------------------------------|
| Machine Name | pln4 |
| Model | HP Proliant DL580 G7 |
| Memory | 512 GB |
| Processors | (4) Intel(R) Xeon(R) CPU X7542 @ 2.67GHz (24 cores) |
| Diskspace | (8) 146.8 GB 15k SCSI |
| Interfaces | (2) 10/1000 |
| FCP HBA | 4/8Gb/s Fibre Channel (configured 8Gb/s) x8 |
| OS | SUSE Linux Enterprise Server 11 (x86_64) - Service Pack 1 |
| Kernel | 2.6.32.12-0.7-default |

| Description | Browse Client Server Hardware Specifications |
|--------------------|---------------------------------------------------------------|
| Machine Name | plsbep2f |
| Model | HP Proliant DL380p G8 |
| Memory | 384 GB |
| Processors | (2) Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz (16 cores total) |
| Diskspace | (16) 300 GB 15k SAS |
| Interfaces | (2) 10/1000 |
| OS | SUSE Linux Enterprise Server 11 (x86_64) Patch Level 3 |
| Kernel | 3.0.101-0.8-default |

| Description | Browse Client Server Hardware Specifications |
|--------------|---------------------------------------------------------------|
| Machine Name | plsbeb2g |
| Model | HP Proliant DL380p G8 |
| Memory | 384 GB |
| Processors | (2) Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz (16 cores total) |
| Diskspace | (16) 300 GB 15k SAS |
| Interfaces | (2) 10/1000 |
| OS | SUSE Linux Enterprise Server 11 (x86_64) Patch Level 3 |
| Kernel | 3.0.101-0.8-default |

CONFIGURATION

Following the best practices, we incorporated the NOOP I/O scheduler, 4k LUN alignment and multipathing.

The IO scheduler was set globally as a menu.lst boot parameter.

The 4k LUN alignment was exercised by using a 4k sector size when creating the XFS file system. This was confirmed by using the vpartial utility.

- **MULTIPATH**

The following multipath.conf was used for the multipath tests. The example templates can be found in /usr/share/doc/packages/multipath-tools. The device configuration settings included in this multipath.conf are recommended in Pure Storage's User Reference Guide. Multipathing was configured in active/active mode.

```
blacklist_exceptions {
    device {
        vendor "PURE"
    }
}

## Use user friendly names, instead of using WWIDs as names.
defaults {
    user_friendly_names yes
    max_fds max
    flush_on_last_del yes
    queue_without_daemon no
}

blacklist {
    wwid 3600508b1001c52e2d818a1196dale48f
    devnode "^hd[a-z]"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    #devnode "^(ram|raw|loop|fd|md|sr|scd|st)[0-9]*"
}

blacklist {
    device {
        vendor "HP"
    }
}

devices {
    device {
        vendor "PURE"
        path_selector "round-robin 0"
        path_grouping_policy multibus
        rr_min_io 1
        path_checker tur
    }
}
```

BENCHMARKS

- **BRANCHSUBMIT**

The Perforce Server (P4D) synchronizes access to its metadata through the use of file locks on the db.* tables. For Perforce tasks that only need to read portions of the metadata, P4D takes read locks on only those db.* tables containing metadata needed by the task. If a task needs to update metadata within a db.* table, P4D takes a write lock on the db.* table. A read lock on a db.* table can be shared with other read locks on the same db.* table, but a write lock on a db.* table is exclusive of all other locks on that db.* table. In general, P4D minimizes the duration that a write lock is held on a db.* table.

One notable exception that can result in P4D holding a write lock on a db.* table for an extended duration is the commit portion of a large changelist submission's dm-CommitSubmit phase. Since the commit portion must be atomic, P4D holds write locks on several important db.* tables for the duration of the commit portion of a changelist submission's dm-CommitSubmit phase. The write locks held block all other tasks that need access to the same tables. It is important that the commit portion of a changelist submission's dm-CommitSubmit phase execute as quickly as possible so that the write locks are released, making the db.* tables available for access to the waiting tasks.

- **BRANCHSUBMIT SUMMARY**

The branchsubmit benchmark results are listed in the following table. The benchmark scenario was run twice, once with 70k files, and once with 700k files per request. The Fiber Channel Protocol (FCP) was used on the Pure Storage FA-420.

While all branchsubmit benchmark statistics are meaningful, the **Commit Rate** and **Elapsed Time** results are the most revealing. The **Commit Rate** (higher results are better) is calculated by the number of files submitted divided by the amount of time a write lock is held on the db.integed table. **Elapsed Time** (lower results are better) is the total amount of time the submit operation takes to complete. Specifically, it is the amount of time the forked child process exists during the submit operation.

| Storage | Run | Commit Rate | Elapsed Time | Compute Phase | Exiting Time | Commit Duration |
|------------------|----------------|--------------------|---------------------|----------------------|---------------------|------------------------|
| Pure Storage FCP | 1 - 70k files | 15024 f/s | 8 sec. | 5605 ms. | 1 sec. | 4659 ms |
| | 2 - 70k files | 21387 f/s | 7 sec. | 5429 ms. | 2 sec. | 3273 ms |
| | 1 - 700k files | 12130 f/s | 80 sec. | 52423 ms. | 3 sec. | 57704 ms |
| | 2 - 700k files | 20037 f/s | 62 sec. | 50678 ms. | 8 sec. | 34934 ms |

- **BROWSE.**

The Browse Benchmark involves a single P4D server and multiple browsechild client machines. Each browsechild instance launched places a load on the server by executing commands that simulate the operational characteristics of the Perforce P4V client. Depending upon the configuration, this test can be CPU and network intensive. Varying the settings for the benchmark configuration can provide information that gauge how well a computer handles a particular load. The browse benchmark focuses on “read” performance.

- **BROWSE SUMMARY**

The results of this benchmark are presented below. These results are meaningful when used to compare against another device’s performance in this test.

The configuration used is:

- FA-420 Pure Storage FlashArray
- FCP (Fibre Channel Protocol)
- Multipathing

| Device/Scheduler/Protocol | 64 children * 2 | 128 children * 2 |
|---------------------------|-----------------|------------------|
| FA-420/noop/FCP | 298 seconds | 668 seconds |

- **PURE STORAGE DATA REDUCTION**

Using Pure Storage CloudAssist, Pure Storage captured a view of the FlashArray after the end of the benchmark testing. The capture revealed a data reduction ratio of 5.2:1.

