PURESTORAGE®
Hello Possible™

MINIGUIDE

ENTERPRISE DBAS ON FLASH: REAL STORIES

Performance Gains, Surprising Survival of an Array-Killing Scenario & Post-Migration DBA Life.

Real Stories: A Big Solutions Group Report

TABLE OF CONTENTS

1 Multi-Vender POC, 1000 DBs and a "Ridiculous" Disk-Array Killing Scenario but "We Couldn't Find the Breaking Point for Pure Storage"
NABEEL SAYEGH

3 Oracle Database Performance Problems Made Quarter-End for DBAs
SOMU RAJARATHINAM

5 Why Enterprise DBA Life With Pure Storage is Technically Easier
CHRIS HINSON

6 Performance Problems with the BI Applications Needed A Robust Solution
CHRIS HINSON

Multi-vendor POC, 1000 DBs and disk-array killing scenario—

STILL CAN'T FIND THE BREAKING POINT FOR PURE STORAGE.

NABEEL SAYEGH

Nabeel Sayegh tested database performance using multiple flash vendors and deployed Pure Storage at his midsize software company.

There was no performance reduction as we loaded up the array. They will tell you to stay below 80%—I have gone over on some occasions, and still, no application layer impact.

We spent about 3-4 months in POC phases with all-flash storage vendors, putting actual production workloads on them for elongated periods of time.

- We not only measured de-dupe/compression rates, but more importantly, the consistency and stability of I/O, capacity utilization and throughput.
- Testing was done while running full re-indexes on over 1000 DB's across 25 SQL (2008 R2) serves, 100 RDS (terminal servers doing a boot storm) and running a combination of Veeam and Avamar backups on about 30 random server types.
- This is NOT something we run during normal production hours, but a ridiculous scenario I came up with that would traditionally kill any disk array.
- In any case, my max IOPS metric was under 197,000 IOPS at 2.9ms latency and 4.5GB/s of throughput.
- We wanted to find the breaking point—we could not find one with Pure Storage.

We have been using Pure Storage for over a year now. We sized our initial production gear based on what we expected from the POC and we were dead on.

We've had positive impacts to our database environments across the board when we moved to Pure Storage. Application response time significantly improved. We proved it from our data cube, but you could also 'feel' it moving around in our applications.

At peak production hours, I have seen these arrays handle near 40,000 IOPS at just over 1ms response (at 70% capacity utilization). Enough said.

The arrays we displaced were EMC Symmetrix VMAX 20K, a DMX- and a VNX 5700 w/flash cache and SSD/SAS auto-tier pool. The cost savings of annual maintenance plus not having to pay for 3-phase power in our co-lo facilities nearly paid for the new Pure Storage arrays.

Pure Storage Support is nothing less than stellar and they continue to be to this day.

- They have a rapid response philosophy and the person who you initially talk to (for whatever the issue) is the one who is resolving it.
- No hand offs and no language barriers. Just people who know what they are doing and bring quick resolution (and explanations) to any issue you have. Even if it is a question (and I have plenty of them), then are quick to respond.

"I am convinced Pure Storage Support have our backs and they keep an eye on things—even if I am not It makes me sleep better at night (and yes, gives me more time for Call of Duty!)"

NABEEL SAYEGH

JUSTIFYING THE COST OF FLASH STORAGE

CHRIS HINSON

We budget about a million dollars a year for capacity increases right now, across the enterprise, for all tiers of storage. In our case, we bought Pure Storage instead of adding capacity to our existing arrays.

After factoring in deduplication and compression, we've found the cost per GB for Pure Storage is cheap enough that we're now replacing Tier 1 and Tier 2 storage with it when those things reach end-of-life. Performance is just a bonus in that equation.

It's not cheaper than the bargain basement disk storage, but it's cheaper than many other things per GB.

For companies that spend $200k or more when they buy storage, Pure Storage should probably be considered. Even with a smaller budget, it should be a consideration just on a cost-per-GB basis.

Oracle Database Performance Problems Made Quarter-End a **WHITE-KNUCKLE EXPERIENCE** at this Enterprise.

SOMU RAJARATHINAM
Oracle DBA Manager and Solutions Architect,
Fortune 1000 company

Before switching to Pure Storage, DBAs used to spend a lot of time on really boring things, like coordinating meetings with the Storage team in planning for striping and mirroring disks and coming up with the perfect LUN size.

Now, with Pure Storage, DBA work is lot more interesting.

We began looking for a solution when our Oracle Data Warehouse database was about 5 terabytes; growing 10 percent a year. We would get email from the senior director about the eroding performance; this was uncomfortable for everyone.

Every quarter it was riskier, we would be bracing for the month end, holding planning meetings with the team to avoid escalations.

SERIOUS PROBLEMS

- 5 terabyte Oracle Data Warehouse plus additional 1 TB database - Refresh took 1 week.
- Backup—over 16 hours.
- Tight storage meant constant firefighting.
- Performance struggled under intense usage.

OUR ATTEMPTED SHORT-TERM FIXES

- Move the database around to reduce the load.
- Housekeeping: cleaning up archived data & old data .
- We talked to legal about retention policy—we had never purged in 7 years.
- Tuning: partitioning a table, standard performance tuning—but we hit limits to what storage can achieve.

Achieving better performance on the database was the key driver that enabled us to justify considering flash storage

For our first POC test, we used a SQL Server database and an Oracle database. Pure Storage came and set up the array, and we were up and running in an hour. We started our test in less than half a day. (This might have taken us a couple of days with another vendor).

"With Pure Storage, our DBAs feel they can get their weekend back to relax and are not stressed out about possible performance issues."

SOMU RAJARATHINAM

KEY FINDINGS AND CAPABILITIES THAT WON US OVER

- Data reduction ratio: 100 GB database was reduced to 33GB.
- Full backups came down from 12-16 hours to 4-6 hours; incremental backups came down from 6-8 hours to 2 hours.
- Snapshotting for backups and replications exceeded expectations.
- This increased our options to consolidate and improve performance significantly.
- New environment can be extended for virtualization (the reduction ratio is better).
- Future-proof ourselves to prevent the same problems cropping up again.

The single Pure Storage POC delivered everything we were looking for. Snapshotting made a big difference to our situation, so much so that we didn't end up testing other systems.

SAVINGS FROM THE SWITCH

- Reduced overall storage footprint; we replaced 3 different storage systems from 3 different vendors, (EMC, NetApp & HP 3Par).
- Improved performance.
- Reduced data center expenses for power, cooling, and rackspace.
- Less staff time spent managing storage.

Before switching to Pure Storage, DBAs used to spend a lot of time on these really boring things, like coordinating meetings with the Storage team in planning for striping and mirroring disks and coming up with the perfect LUN size. Now, with Pure Storage, DBA work is lot more interesting.

HERE ARE SOME EXAMPLES

- Start planning consolidation of various other databases that are in different storage systems.
- Enable database cloning for various other databases—which would have been a luxury earlier.
- Look for automation opportunities to clone the environments to improve overall productivity.
- Focus on keeping up with latest Database versions rather than firefighting performance and space management.

WHY ENTERPRISE DBA LIFE WITH PURE STORAGE IS TECHNICALLY EASIER

CHRIS HINSON
Enterprise DBA

DBAs are usually separated from the body of work required of the storage admins, but I can give you a good example of something I don't have to think about anymore on a technical level.

With spinning disk, RAID levels matter for different types of files. Ideally, you'll use RAID 10 for Logging and transactional data, but maybe RAID 5 for historical or archive data. That's a generalization, but you get the idea.

With Pure Storage, I don't have to worry about RAID levels for Flash. I just ask for space—which drives goes where—and don't have to get more specific than that.

The Pure Storage software intelligently makes all these decisions based on what's happening: The RAID decisions, the number of drives—you don't mess with it.

With storage and performance, there are so many variables: the hardware itself, if you have more drives, if you need more bandwidth, if you have fewer dries, how is it connected, who's configuring it.

And if you have 100 disks allocated in a traditional array, RAID types, drives training, replication, different storage volumes on different servers, many configuration options—saying "this is the right way" or "this is the wrong way"—it's such black magic.

There is so much configuration required to get it going and it gets way down into the weeds on how it works technically, buffer memory, read and write mix, how many disks allocated to what.

There's a lot more involved in the "forklift upgrade" and if you take the out of the box storage configuration on a lot of flash arrays it's probably not going to be the best for the situation.

Performance Problems with the BI Application **NEEDED A ROBUST SOLUTION**

CHRIS HINSON
Senior Database Administrator, Fortune 500 company.

He's worked on several different DBMS (Oracle, MS SQL and Informix) and storage platforms.

"Moving onto flash storage has changed the way some of my biggest DBA tasks work and given me new options on how to approach some big challenges."

CHRIS HINSON

We have generally good infrastructure and database systems performance. Our main problem was with our BI system, which doubled in size every year.

- The BI manager was getting pressure from the business users (Sales, Marketing, Pricing, Operations) to get reports faster and earlier in the day.
- Users also wanted more functionality and new datasets—so more data growth.
- That pressure was passed on to the infrastructure and database teams, as well as the BI development team.
- BI is married to a reporting platform (Cognos) and ETL tool (DataStage), and needed better performance from the database and app servers.
- Unsurprisingly, these demands weren't well qualified—they just wanted more/better/faster.

As we pursued hardware solutions, Pure Storage was our 3rd POC. We loaded all of our BI databases on Pure Storage and hammered it with backups and restores, full DB index rebuilds, compressed tables, anything we could think of. We even broke one of the controllers, but it didn't cause an interruption in our testing.

- Even though we were using compression on the DB side, we still got 4-5:1 compression on the array, which means a really nice cost-per-gig proposition.
- Performance-wise, we were consistently running a full gigabyte per second of throughput from the single BI machine—and running similar numbers from other machines, simultaneously.

HEAVY ETL LOAD IS GONE, BI USER SATISFACTION IS UP

A few years ago, we had a single database instance that was our BI reporting and ETL server. It was a terrible system.

- Users were allowed to run reports at any time, even while the ETL was running (because it took all day).
- Report had inconsistent results as tables were loaded.
- And, the reports would block ETL activities.
- We eventually split that server into a reporting instance and an ETL instance.
- We backed up the ETL instance daily and restored to the Reporting instance multiple times a day using differentials.
- This was a huge hit to the disk-based storage and made performance poor for every other app on that same SAN.

This situation was one of the primary reasons we bought Pure Storage.

- The snap/copy technology makes our BI system available for reporting more quickly—and without the backup/restore overhead.
- **We now have a fully automated server job that initiates a snap copy of the 5 TB ETL DB and loads it onto the reporting server in a matter of minutes.**
- Now, reports are available in the early morning, which means users are happy.
- The fact that the backup and restore is eliminated is another huge win for every other app in the company aside from BI, because that storage I/O overhead isn't happening on the SAN anymore so everybody else has more bandwidth.
- Application owners outside of BI have heard about the success and performance with Pure Storage and have come asking to be migrated.

> *"With the changes we've made, I have more freedom to enjoy my time off, since I don't need to worry about being tethered to my laptop for urgent issues.*
>
> *Saturday, I took my 3-year-old daughter hiking and didn't need to worry about anything."*
>
> **CHRIS HINSON**

Here's another good example of how much easier DBA life is with Pure Storage:

- In my BI app, we just released new code in production this weekend.
- Today, I'm using a copy of Production replicated from Las Vegas to Phoenix to refresh my full volume Dev and QA databases.
- Instead of doing a new backup (5 TB—already compressed—that with really good compression turns into maybe 500-1000 GB) and then taking hours to copy that backup across data centers, I've got a replicated copy that was done over the wire once—months ago—and now only has to replicate changes (FlashRecover Replication).
- Then, I ask for a snap copy of that data and it's loaded on to the Dev and QA servers, instead of restoring.
- The process takes minutes instead of hours and uses very little additional storage space (FlashRecover Snapshots).

As time goes on, we're adding more features and doing more things with Pure Storage. For example, we're doing Virtual Desktop Infrastructure (VDI) now and the dedup/compress for that workload is great (20:1-ish).

THE ORIGIN OF REAL STORIES

Tech pros seek insights and share unvarnished opinions in independent forums all over the web. That's where this Real Stories project & research started. This report is drawn entirely from Pure Storage Real Users' words, observations and experiences. All Stories are used with permission.

ABOUT PURE STORAGE, INC.

Pure Storage accelerates possible, transforming businesses in ways previously unimagined. The company's disruptive, software-driven storage technology combined with a customer-friendly business model drives business and IT transformation for customers through dramatic increases in performance and efficiency at lower costs. Pure Storage FlashArray//m is simpler, faster and more elegant than any other technology in the datacenter. FlashArray//m is ideal for the move toward big data and for performance-intensive workloads such as cloud computing, database systems, desktop virtualization, real-time analytics and server virtualization. With Pure's industry leading NPS score of 79, Pure customers are some of the happiest in the world, and include large and mid-size organizations across a range of industries: cloud-based software and service providers, consumer web, education, energy, financial services, governments, healthcare, manufacturing, media, retail and telecommunications. With Pure Storage, companies push the boundaries of what's possible to become faster, smarter and more innovative.

To learn more, visit http://purefla.sh/Oracle