



MINI GUIDE

Enterprise DBAs On Flash: **REAL STORIES**

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




TABLE OF CONTENTS

03 **CHAPTER 1**
Multi-Vendor POC, 1000 DBs
and a “Ridiculous” Disk-Array
Killing Scenario but “We Couldn’t
Find the Breaking Point for
Pure” [Nabeel Sayegh](#)

06 **CHAPTER 2**
Why Enterprise DBA Life With
Pure is Technically Easier
[Chris Hinson](#)

07 **CHAPTER 3**
Performance Problems with
the BI Applications Needed
A Robust Solution
[Chris Hinson](#)

ONE

Multi-Vender POC, 1000 DBs and a “Ridiculous” Disk-Array Killing Scenario but

“WE COULDN’T FIND THE BREAKING POINT FOR PURE”

“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.”

NABEEL SAYEGH

We spent about 3-4 months in Proof Of Concept 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.

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 very predictable.

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.

“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

We’ve had positive impacts to our database environments across the board when we moved to Pure. 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 arrays.”

—

NABEEL SAYEGH

Pure Storage’s 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.

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

A close-up photograph of a person's hands. One hand is pointing at a document on a wooden table, while the other hand holds a red fountain pen. The document features several blue bar charts and a pie chart. The background is a blurred office setting with wooden chairs and tables.

JUSTIFYING THE COST OF FLASH STORAGE

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 instead of adding capacity to our existing arrays.

After factoring in deduplication and compression, we've found the cost per GB for Pure 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 should probably be considered. Even with a smaller budget, it should be a consideration just on a cost-per-GB basis.

TWO

Why Enterprise DBA Life With Pure is **TECHNICALLY EASIER**

CHRIS HINSON // SENIOR DATABASE ADMINISTRATOR, FORTUNE 1000 COMPANY

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, 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.”

—

CHRIS HINSON

The Pure 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 drives, 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, there are so many configuration options—saying this is the right way this is the wrong way—it's such a 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.

THREE

Performance Problems with the BI Application Needed

A ROBUST SOLUTION

“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 // SENIOR DATABASE ADMINISTRATOR, FORTUNE 1000 COMPANY

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 was our 3rd POC. We loaded all of our BI databases on Pure 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.

- 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 and have come asking to be migrated.

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

- 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 datacenters, I've got a replicated copy that was copied over the wire once (months ago) and now only has to replicate changes (Pure 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 (Pure snap copies).

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

Chris Hinson is Senior Database Administrator at a Fortune 500 company.

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

“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.

So on Saturday, I took my 3-year-old daughter hiking and didn't need to worry about anything.”

CHRIS HINSON //
SENIOR DATABASE ADMINISTRATOR,
FORTUNE 500 COMPANY



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 helps innovators build a better world with data. Pure's data solutions enable SaaS companies, cloud service providers, and enterprise and public sector customers to deliver real-time, secure data to power their mission-critical production, DevOps, and modern analytics environments in a multi-cloud environment. One of the fastest growing enterprise IT companies in history, Pure Storage enables customers to quickly adopt next-generation technologies, including artificial intelligence and machine learning, to help maximize the value of their data for competitive advantage. And with a Satmetrix-certified NPS customer satisfaction score in the top one percent of B2B companies, Pure's everexpanding list of customers are among the happiest in the world.