

DATABASES ON ALL-FLASH ARRAYS

How it changes the life of the DBA

April 2018

Istvan Stahl

Gabor Akots


1. Database refresh challenges
2. Database refresh requirements
3. Available options
4. Traditional Block Storage Arrays vs Pure Storage
5. Benefits of using Pure storage refreshes
6. How it works
7. Lessons learned
8. Q/A

DBA CHALLENGES

- Refresh non-Production environments with Production data fast
- The amount of data to be managed is exploding
- Serve Agile development requirements

DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used

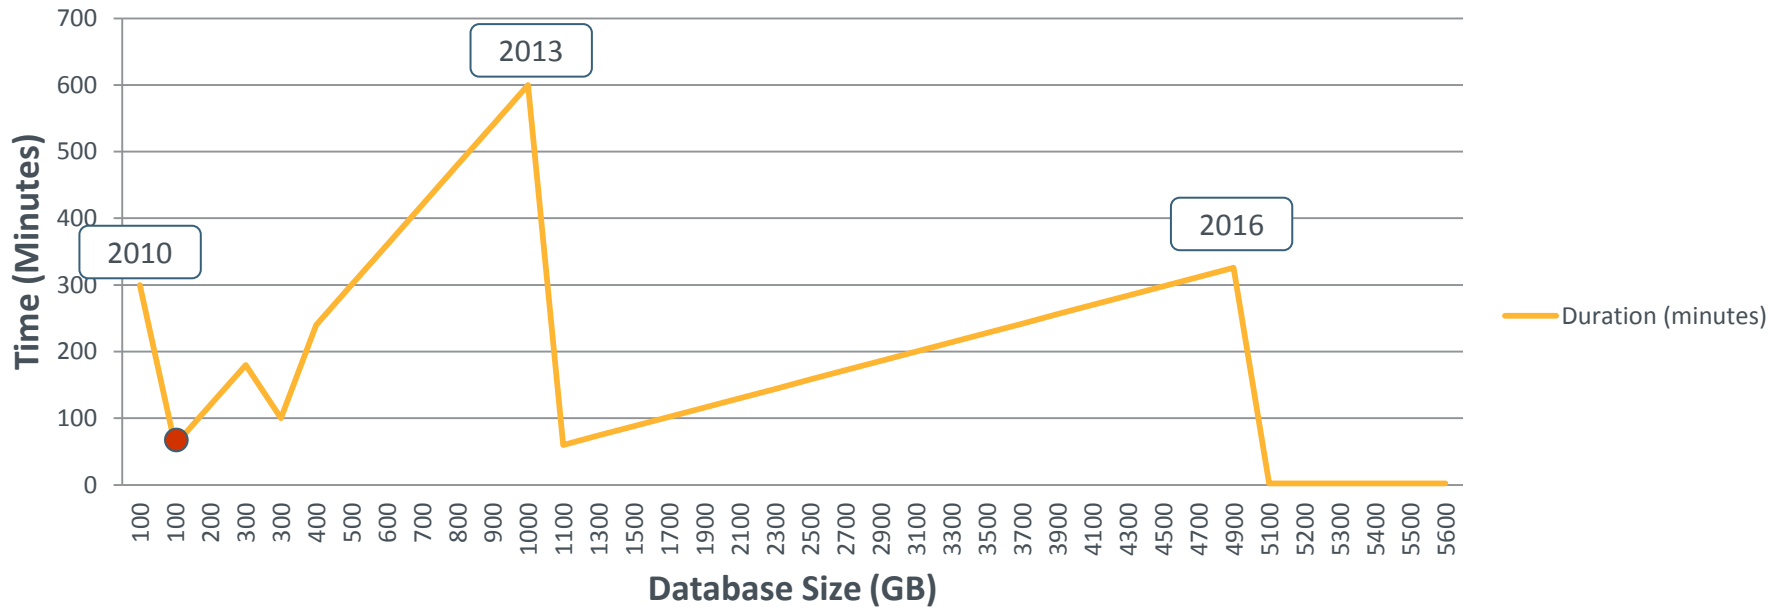
Customer satisfaction

Manual DB refresh using exp/imp



DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used


Customer satisfaction

Self Service Portal leveraging expdp/impdp



DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used


Self Service Portal leveraging expdp/impdp
(getting slower)

Customer satisfaction



DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used

Manual DB refresh using expdp/impdp


Idea do not validate the constraints during import

Result optimizer behaviour changes



DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used

Self Service Portal leveraging expdp/impdp
(this is very slow)


ISSUE: Datapump eliminates row chaining

Customer satisfaction



DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used

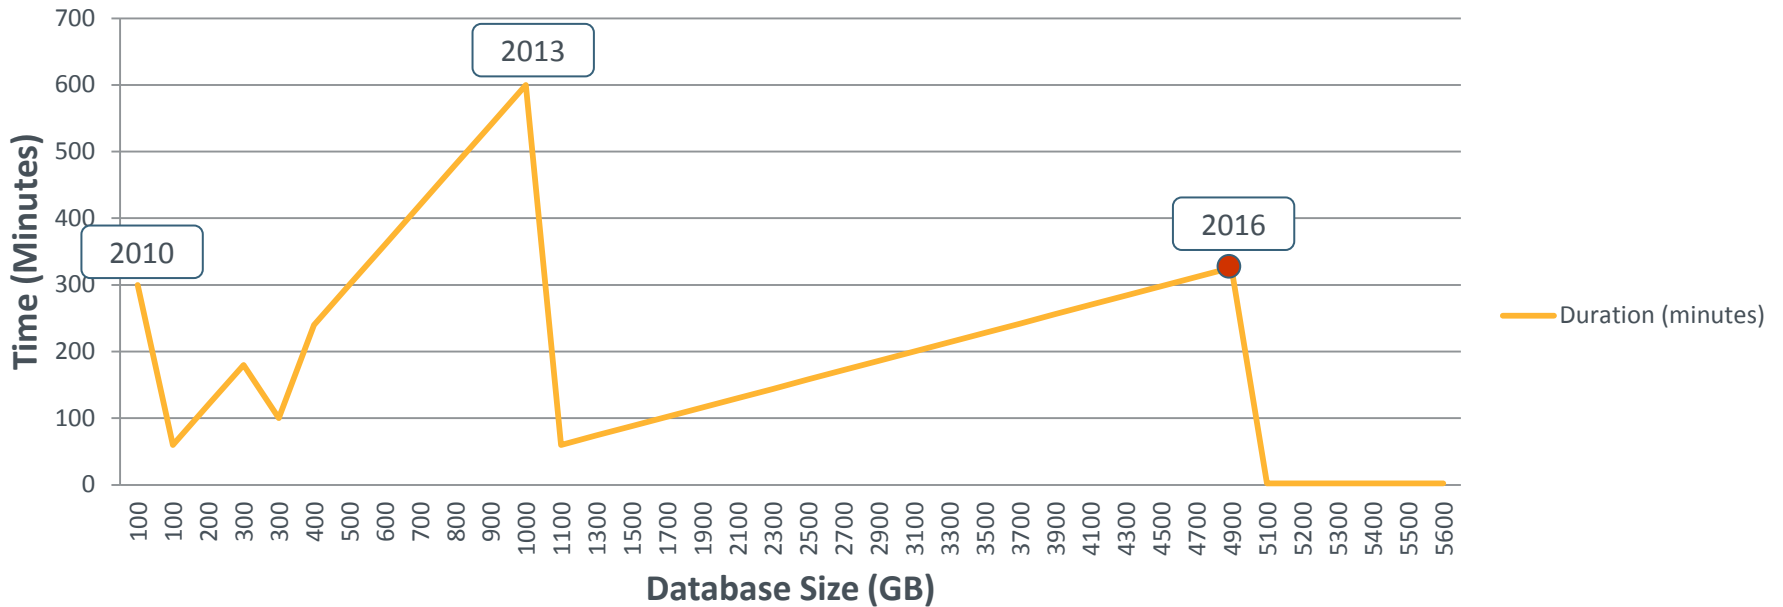
RMAN Restore

Customer satisfaction



DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used


Customer satisfaction

RMAN restore (things are slow again)



DEV/UAT/QA DATABASE REFRESH CHALLENGES

Database refresh speed



Technology used

Customer satisfaction

Introducing All-Flash Storage snapshots



DATABASE REFRESH OPTIONS

- Oracle
 - Multitenant 12.1 Snapshot Clones (Oracle CloudFS aka ACFS)
 - **12.2 New Features:** Hot PDB Clones, Refreshable PDB
 - dNFS Clonedb
 - Enterprise Manager Snap Clone
- Third party



Do it yourself leveraging  **PURE**STORAGE *!!!*



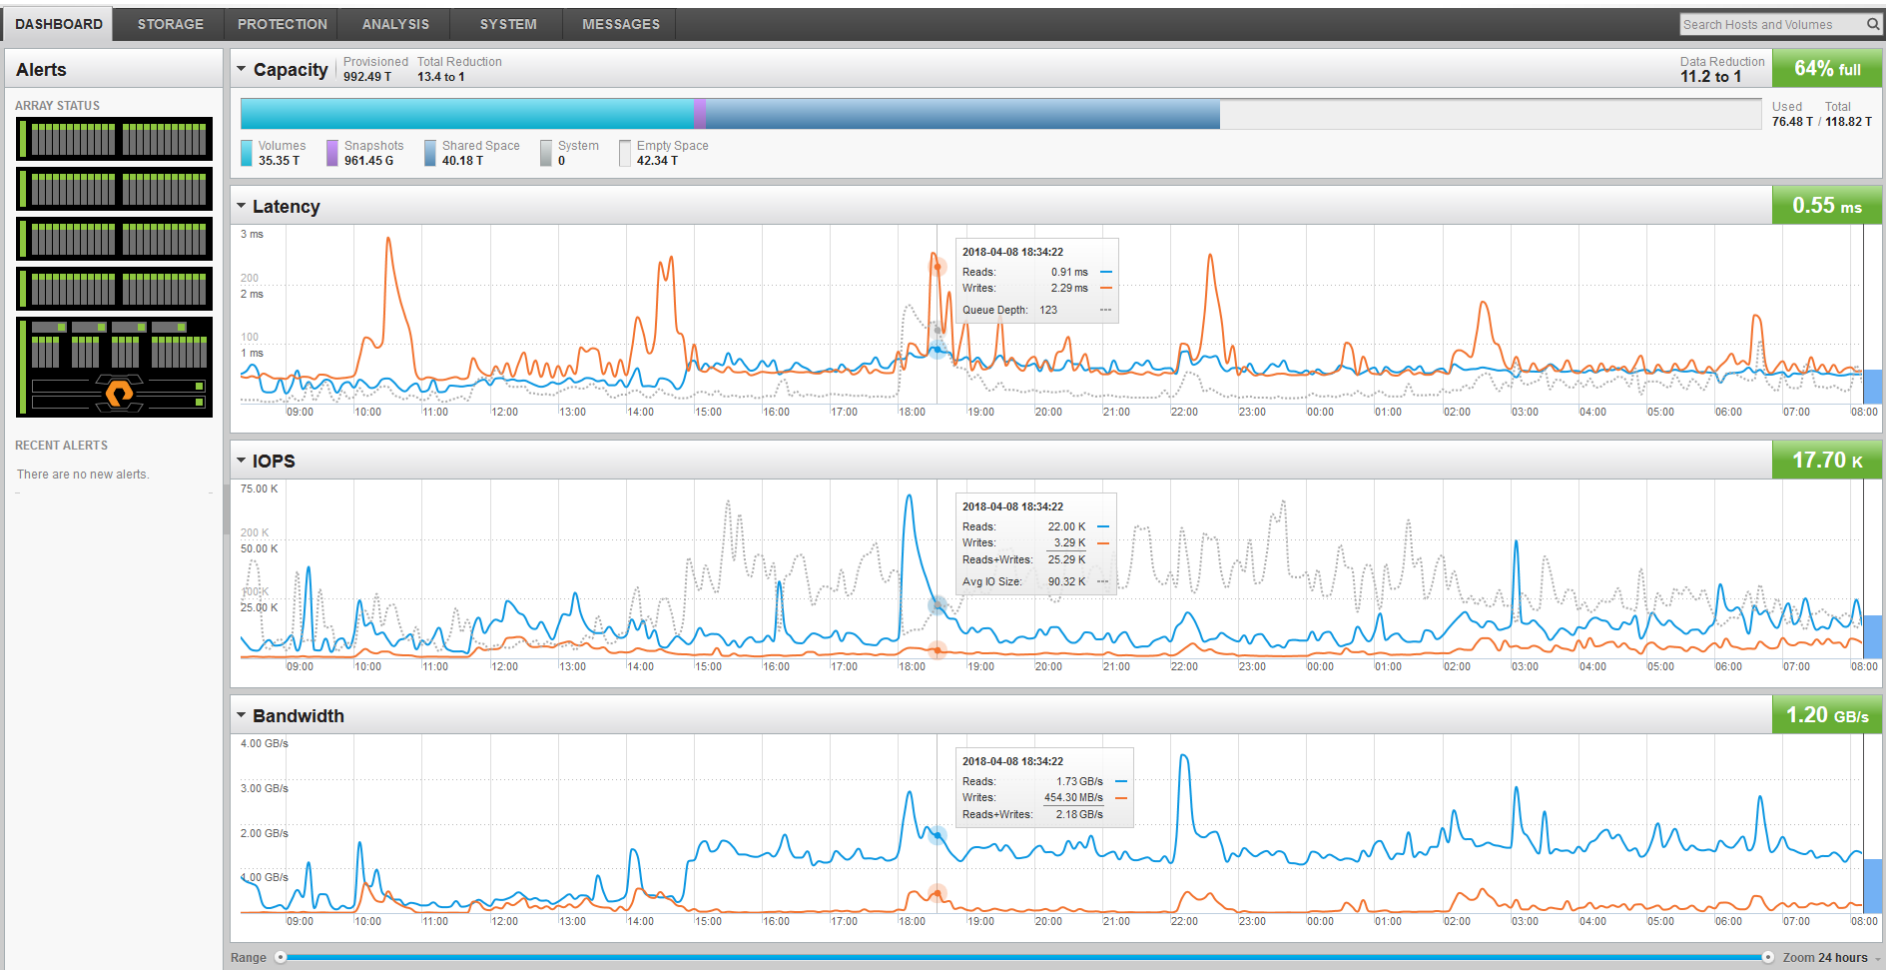
PURE VS TRADITIONAL BLOCK STORAGE ARRAYS

- Intelligent features: Deduplication, compression
- Total capacity of the array
- Complexity, manageability (single lun diskgroups)
- Latency below 1 ms
- Replication: sync replication was not available in Pure when we started
- Upgrade path: Evergreen vs. Buy a new array every 3rd year
- Density
- Power consumption

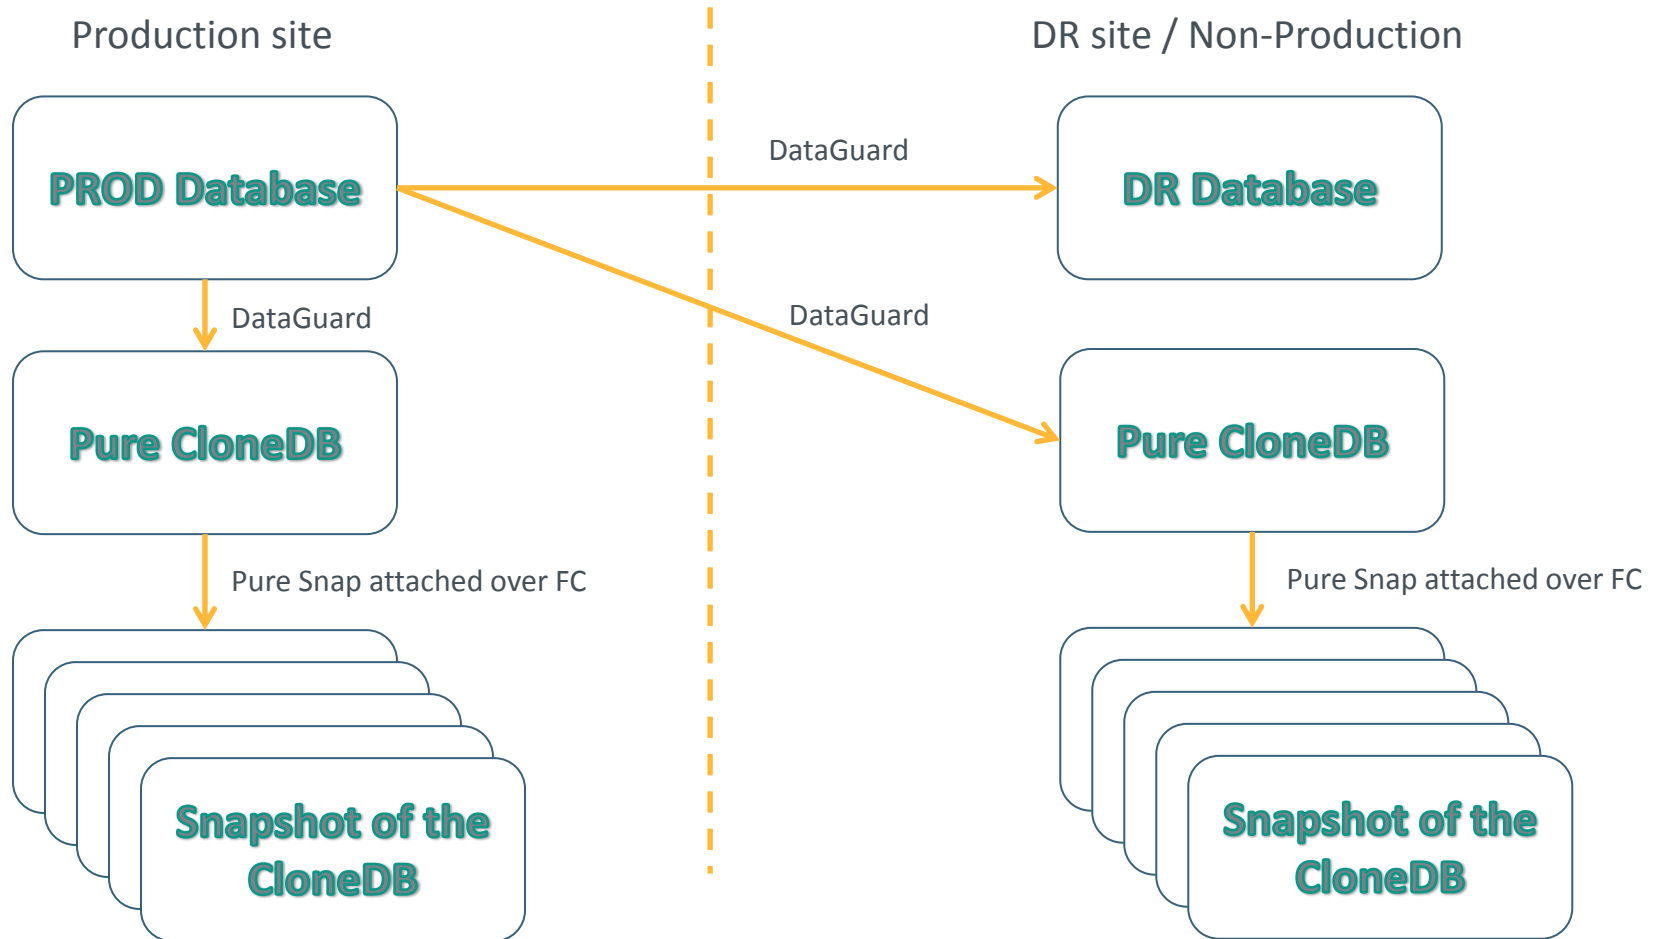
BENEFITS OF USING PURE STORAGE REFRESHES

	Traditional Block Storage	Pure Storage
Size of the Primary database	5 TB	1.2 TB
Number of Dev/QA/UAT environments	40	40
Total space consumption	200 TB	10TB
Total cost of Storage (Capex & Opex)	\$x USD	10%
Data Reduction	1:1	20:1
Environment refresh time (one environment)	6 hours	2 minutes
Annual Business time saving (baseline: Traditional)	0	6000 hours
Annual DBA Ops time saving (baseline: Traditional)	0	2000 hours
Annual data transfer over the network	5 PB	0 PB

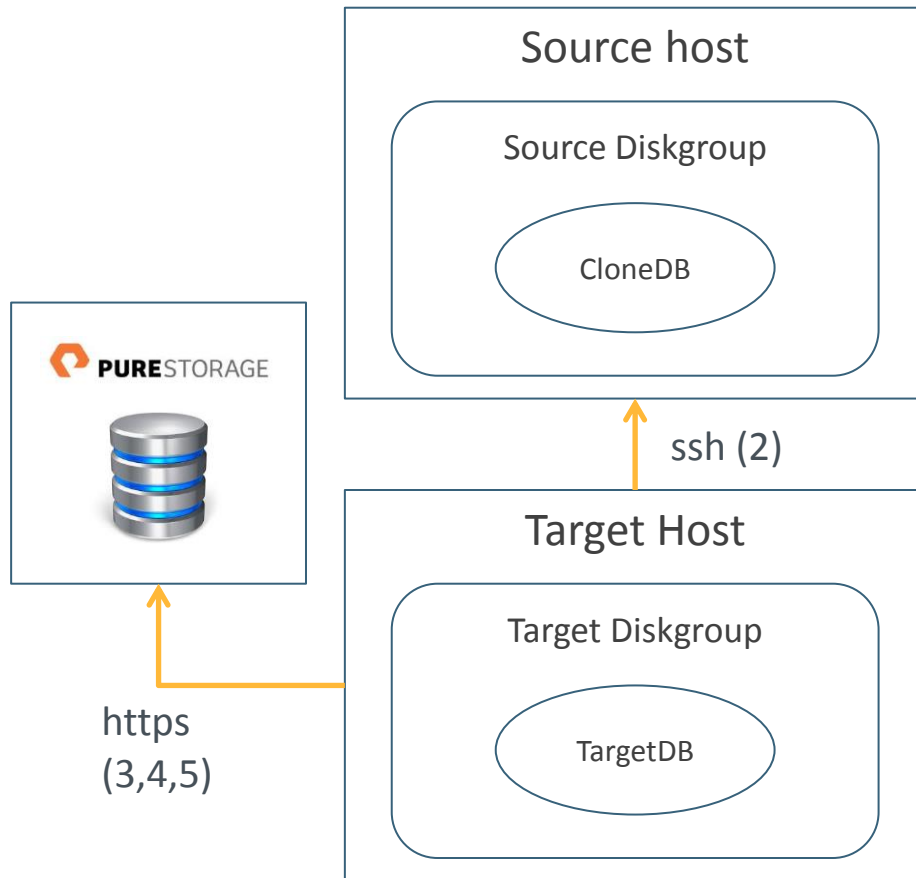
PURE STORAGE WEB CONSOLE



HOW IT WORKS



HOW IT WORKS



1. Cloning process starts on the target Host
2. Get disk details of the Source Diskgroup (ssh)
3. Take a snapshot of the disks on Pure (https) *
4. Copy the snaps to the new target volumes (https)
5. Connect the volumes to the Target Host (https)
6. Rescan the scsi bus on Target Host
7. Setup proper device permissions for ASM
8. Rename Source DG to TargetDG
9. Mount Target Diskgroup
10. Clone CloneDB to TargetDB (TargetDB will be a Primary database)
11. Post refresh actions (eg.: roll forward, database upgrade)

* - we take consistent snaps of a dataguard replica always

LESSONS LEARNED

- Pure does deduplication in multiple rounds
- Databases prior 12.1 support only 2TB disks, even when the ASM layer supports larger disks (chance for data corruption)
- The Application might not benefit from All-Flash
 - When it uses the default pre-fetch size
 - When all data sits in memory and the processing is not I/O heavy
- TDE Encryption kills storage level Compression, Deduplication still helps a lot
- Pure Storage is really simple (even a DBA can understand 😊), fast and its Rest API rocks

Q / A