

NetApp FabricPool

Тиринг нового поколения

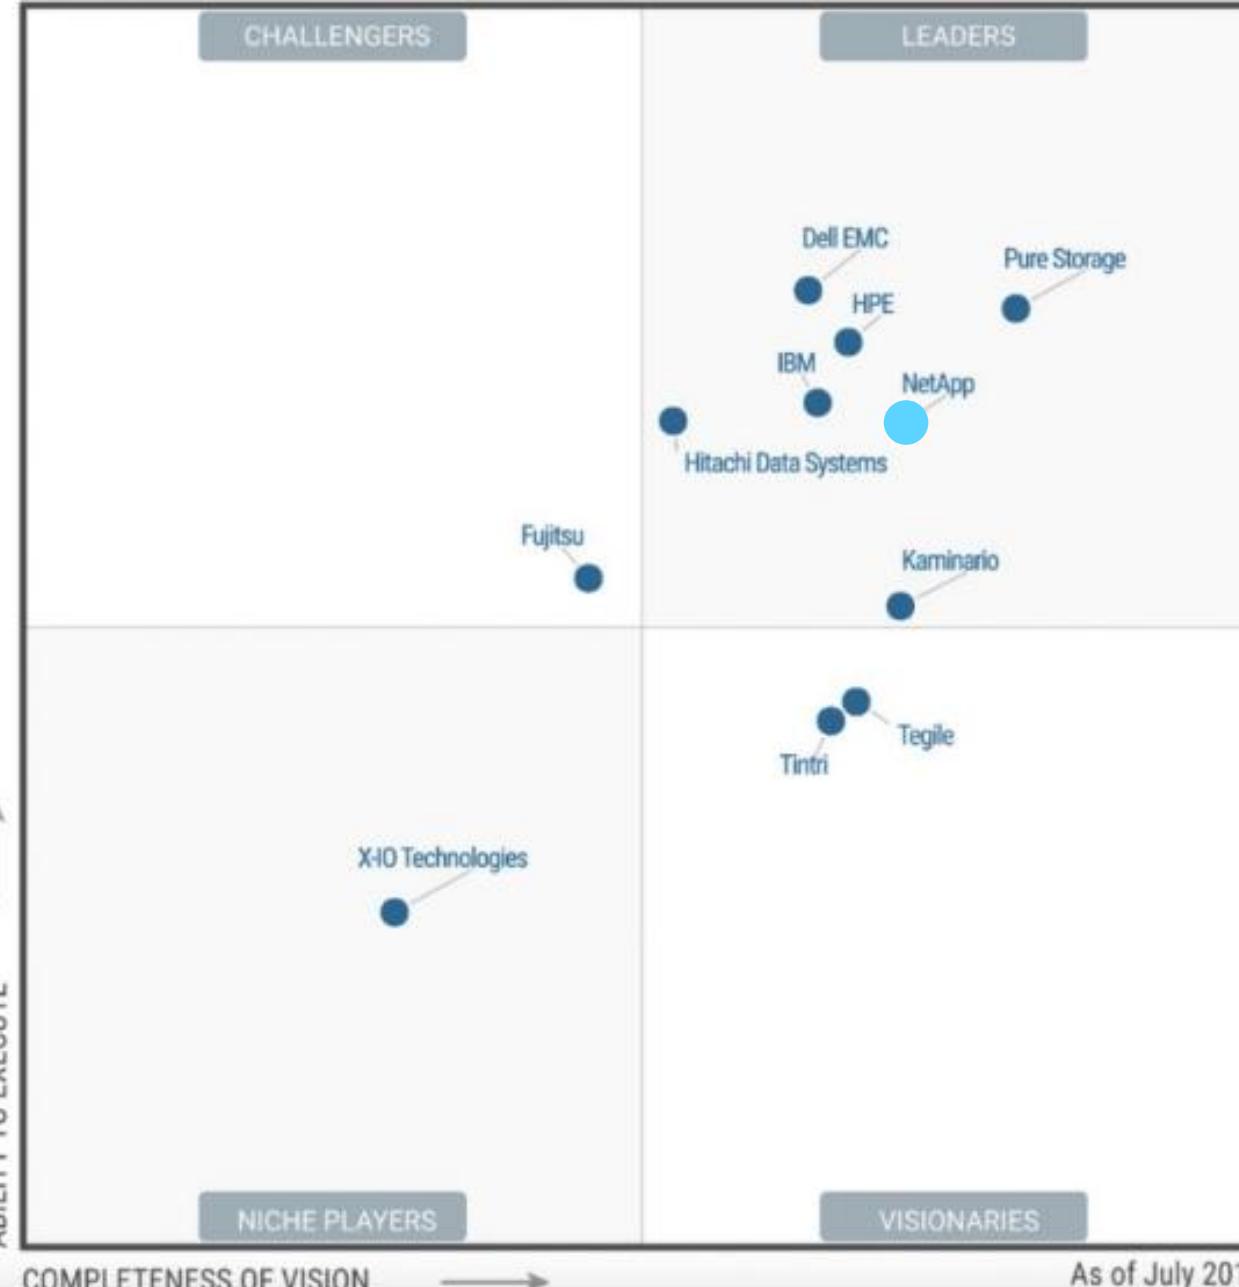
Кириленко Антон

Инженер СХД

Мегатрейд

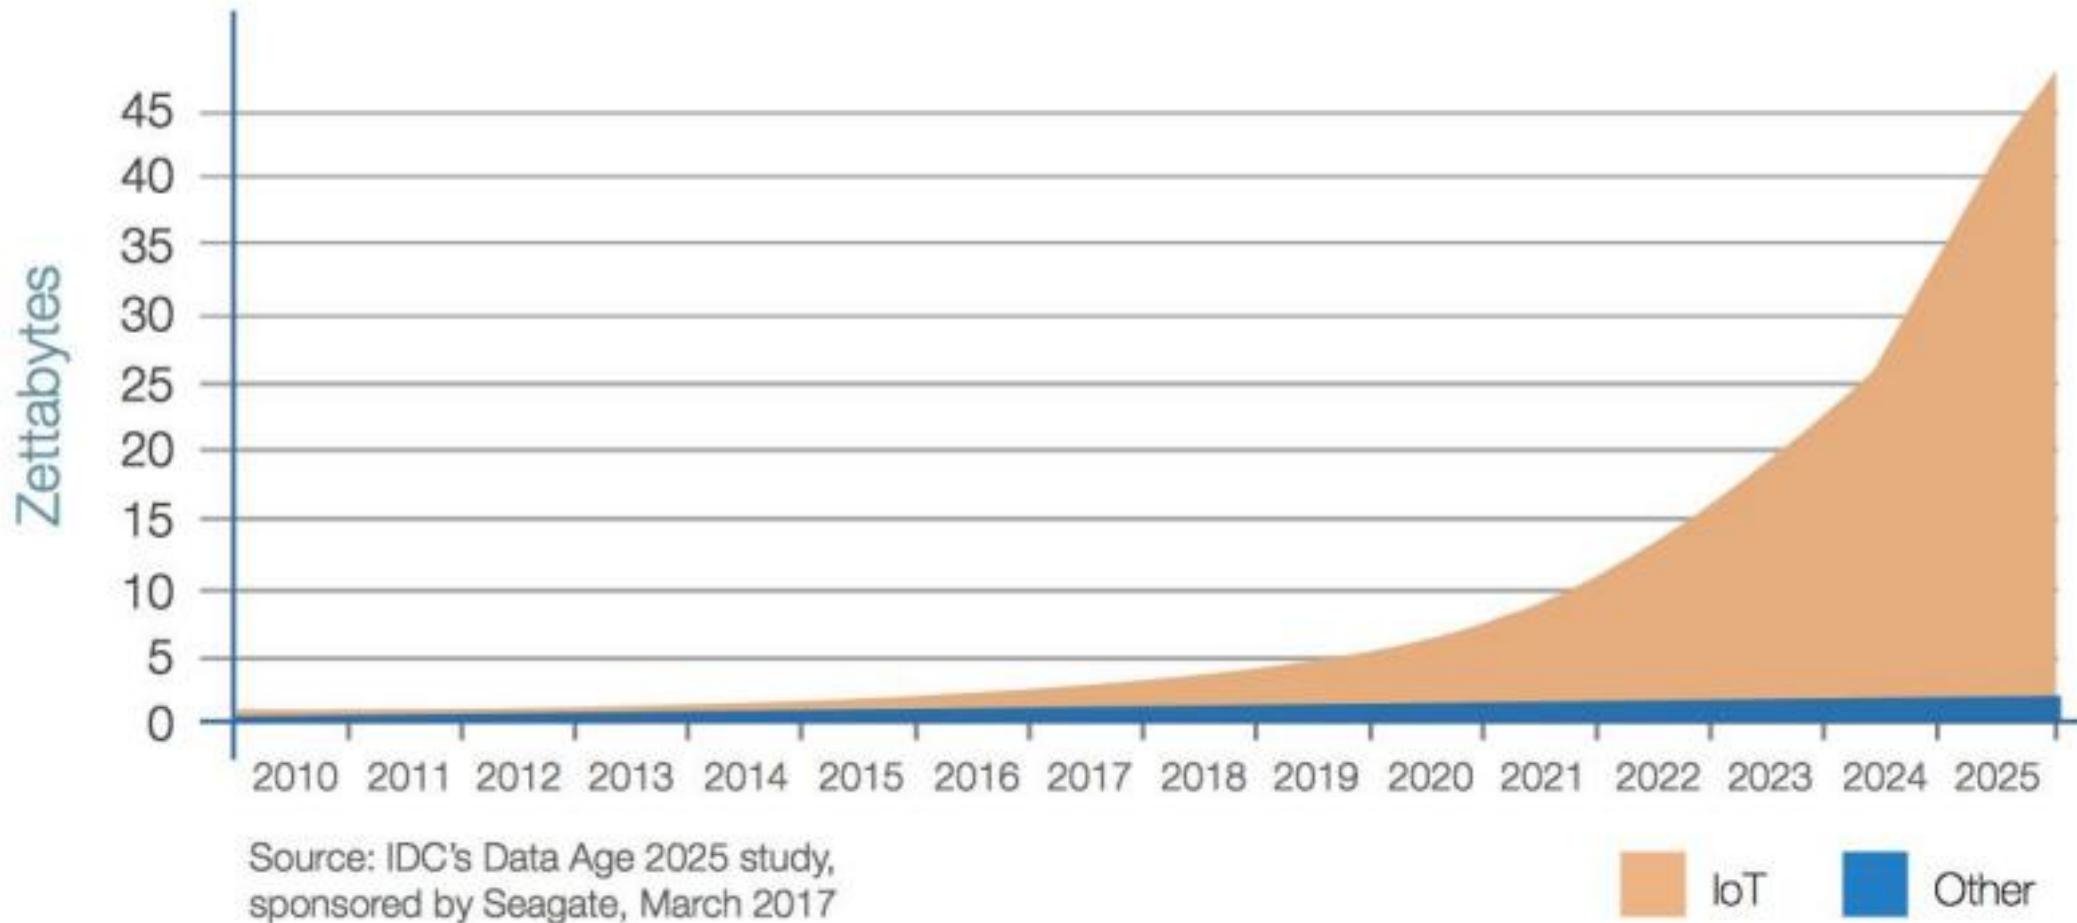


Magic Quadrant for Solid-State Arrays 2017

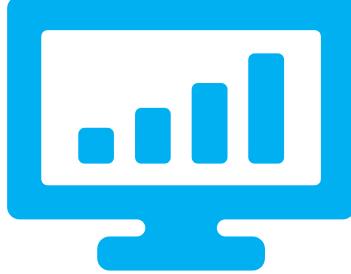


- К 2021 году 50% ЦОД будут использовать твердотельные массивы (SSA) для высокопроизводительных вычислений и больших рабочих нагрузок, по сравнению с менее чем 10% сегодня
- В течение следующих 12 месяцев твердотельные массивы (SSA) улучшатся в производительности в 10 раз и удваиваются по плотности

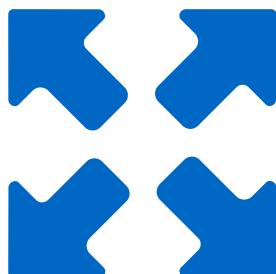
Data Tsunami



Увеличиваем производительность, уменьшаем стоимость



- **Inline adaptive compression**
 - Потоковая компрессия
- **Inline deduplication**
 - Потоковая дедупликация
- **Inline data compaction**
 - Потоковое уплотнение
- **Fabric Pool (Tiering)**
 - Тириング холодных данных



NetApp Storage Efficiency portfolio

AFF ONTAP 9.x

1) Data deduplication

- Inline (**on** by default)

2) Data compression

- Inline (**on** by default)

3) Data compaction

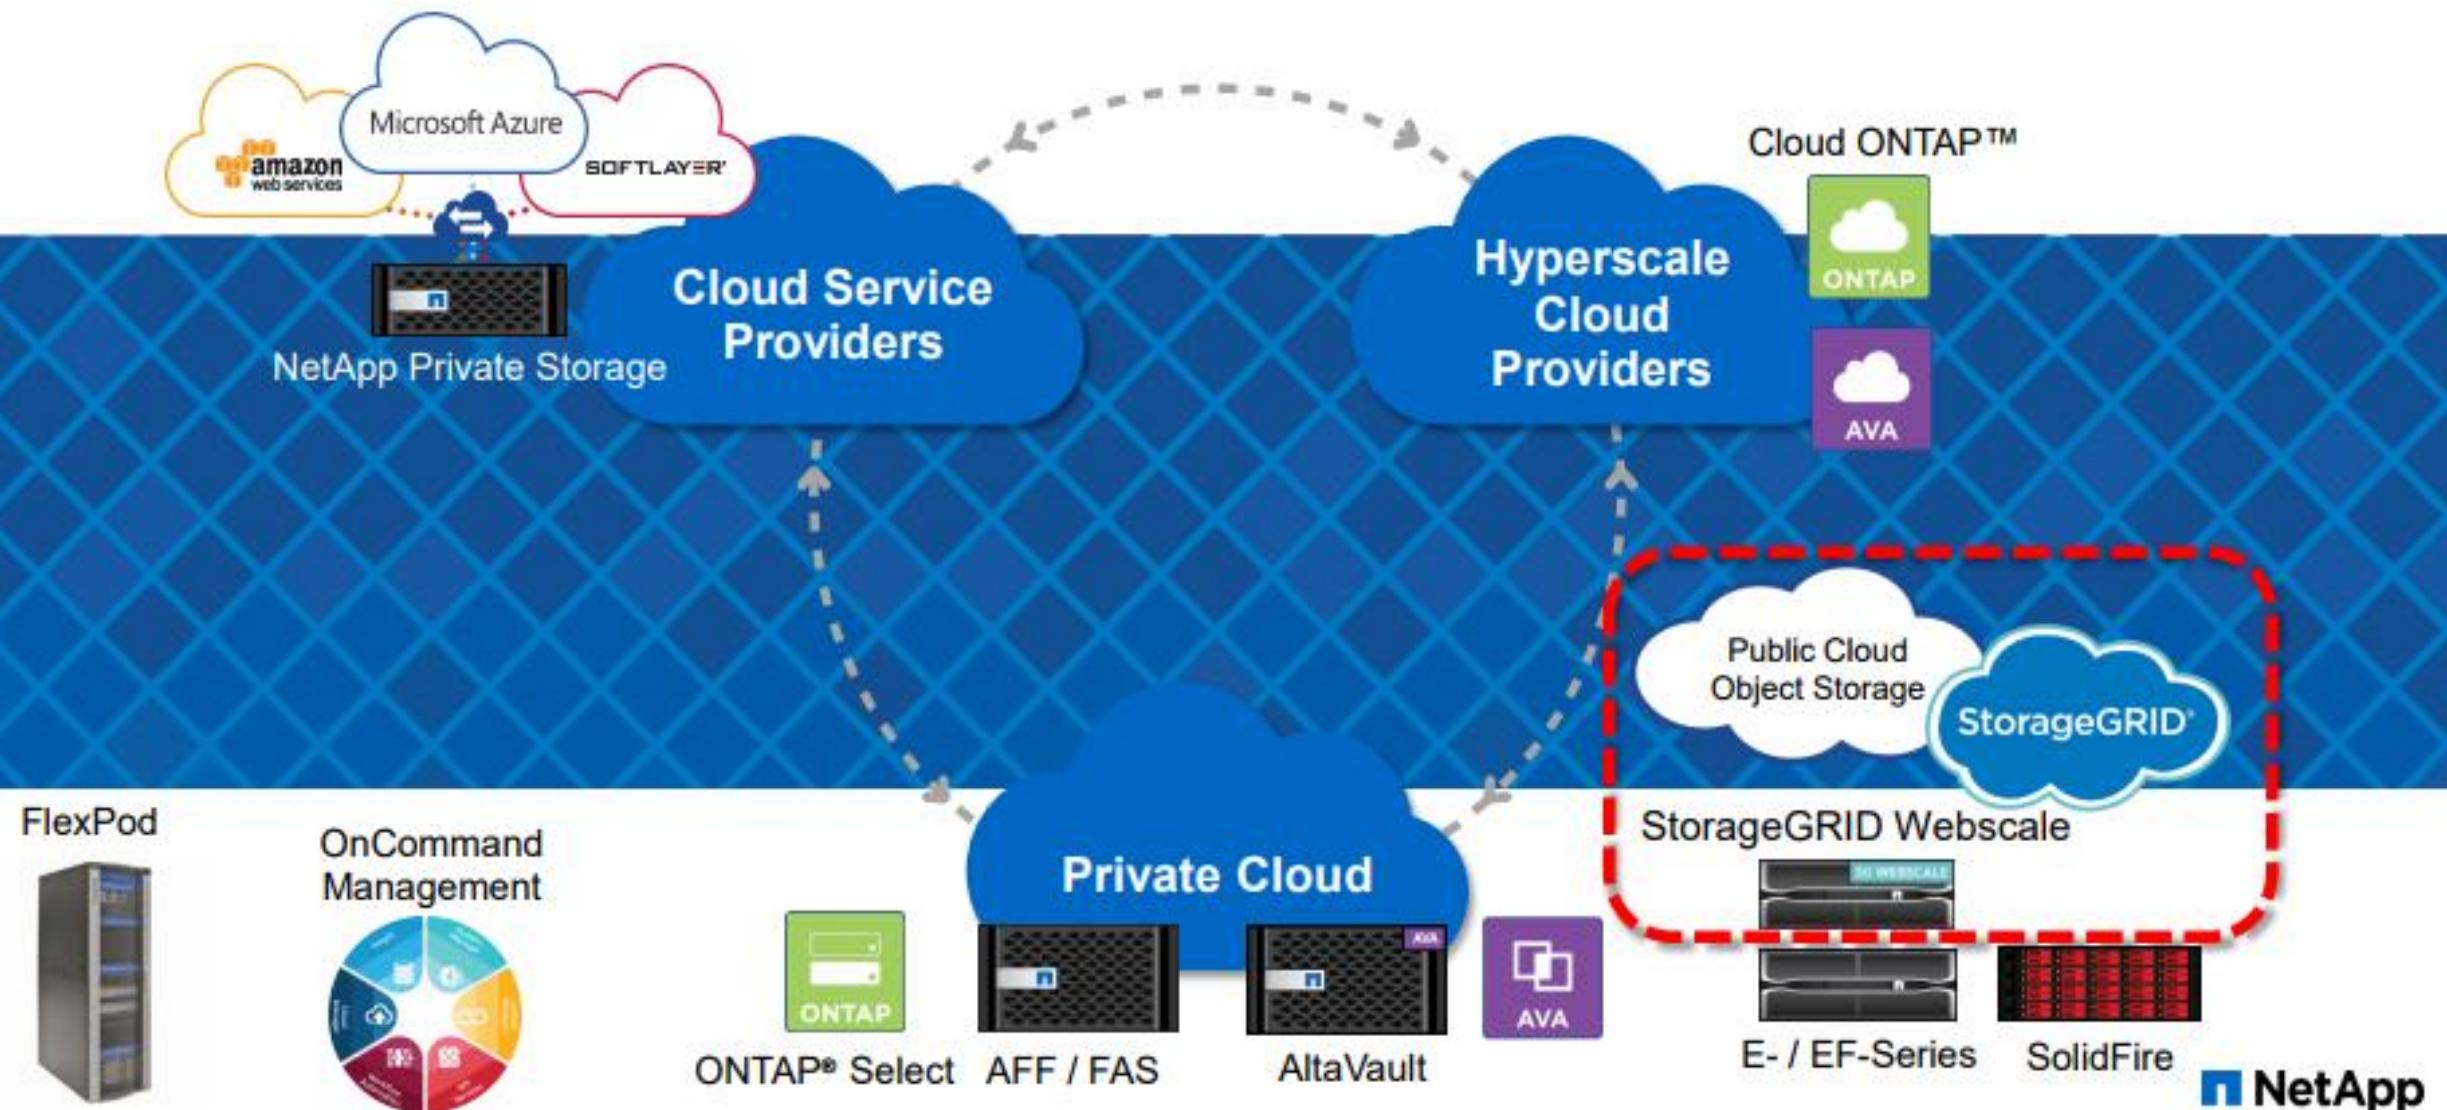
- Inline (**on** by default)

AFF Storage Efficiency Ratios with ONTAP® 9

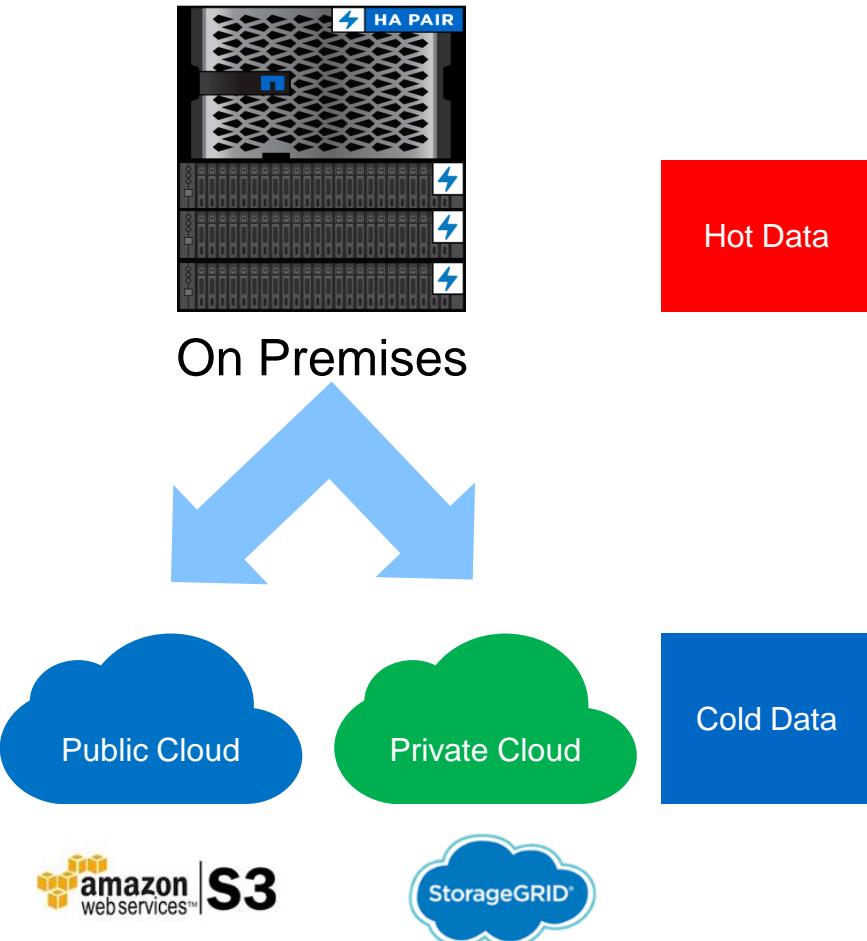
| Workload | Range |
|--|-------------|
| Oracle databases (with no database compression) | 2:1 – 4:1 |
| SQL Server 2014 databases (with no database compression) | 2:1 – 4:1 |
| VMware Horizon full clone desktops (persistent) – NetApp® clones | 5:1 – 10:1 |
| VMware Horizon linked clone desktops (non-persistent) | 3:1 – 5:1 |
| Citrix XenDesktop full clone desktops (persistent) – NetApp clones | 5:1 – 10:1 |
| Citrix XenDesktop MCS desktops (non-persistent) | 3:1 – 5:1 |
| Citrix Provisioning services desktops (non-persistent) | 3.3:1 – 5:1 |
| Virtual Servers (OS and applications) | 2:1 – 4:1 |
| Microsoft Exchange | 1.6:1 |
| Home directories | 1.5:1 – 2:1 |
| Software development | 2:1 – 10:1 |

Что же такое FabricPool?

Data Fabric

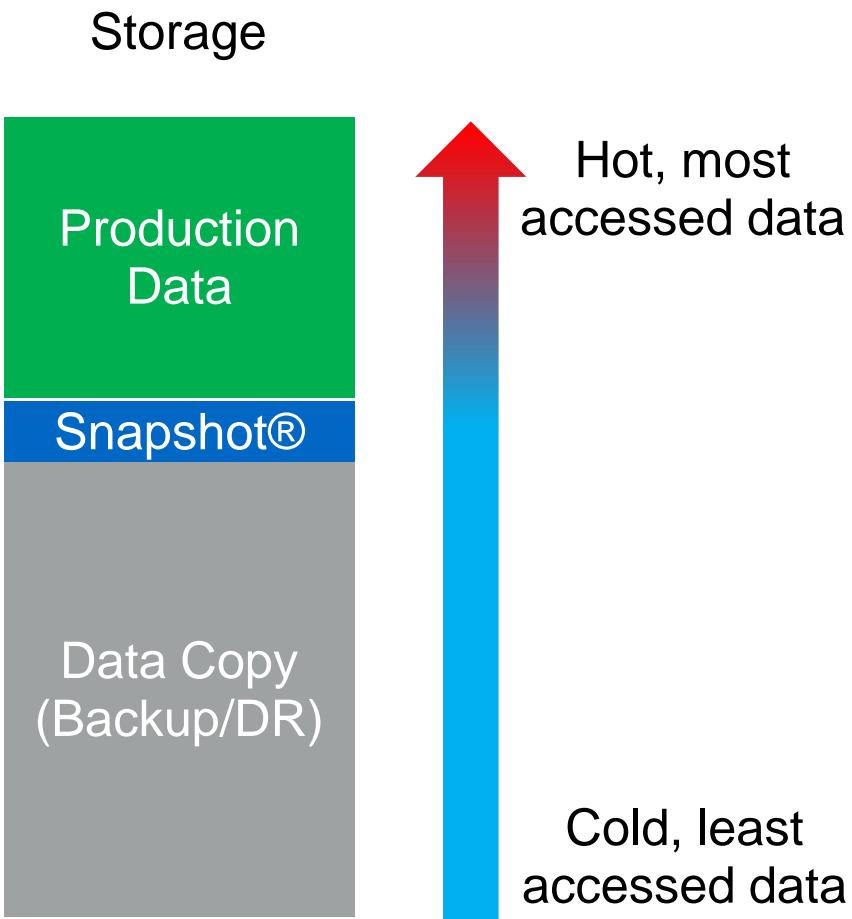


Автоматизируем облачный тиринг



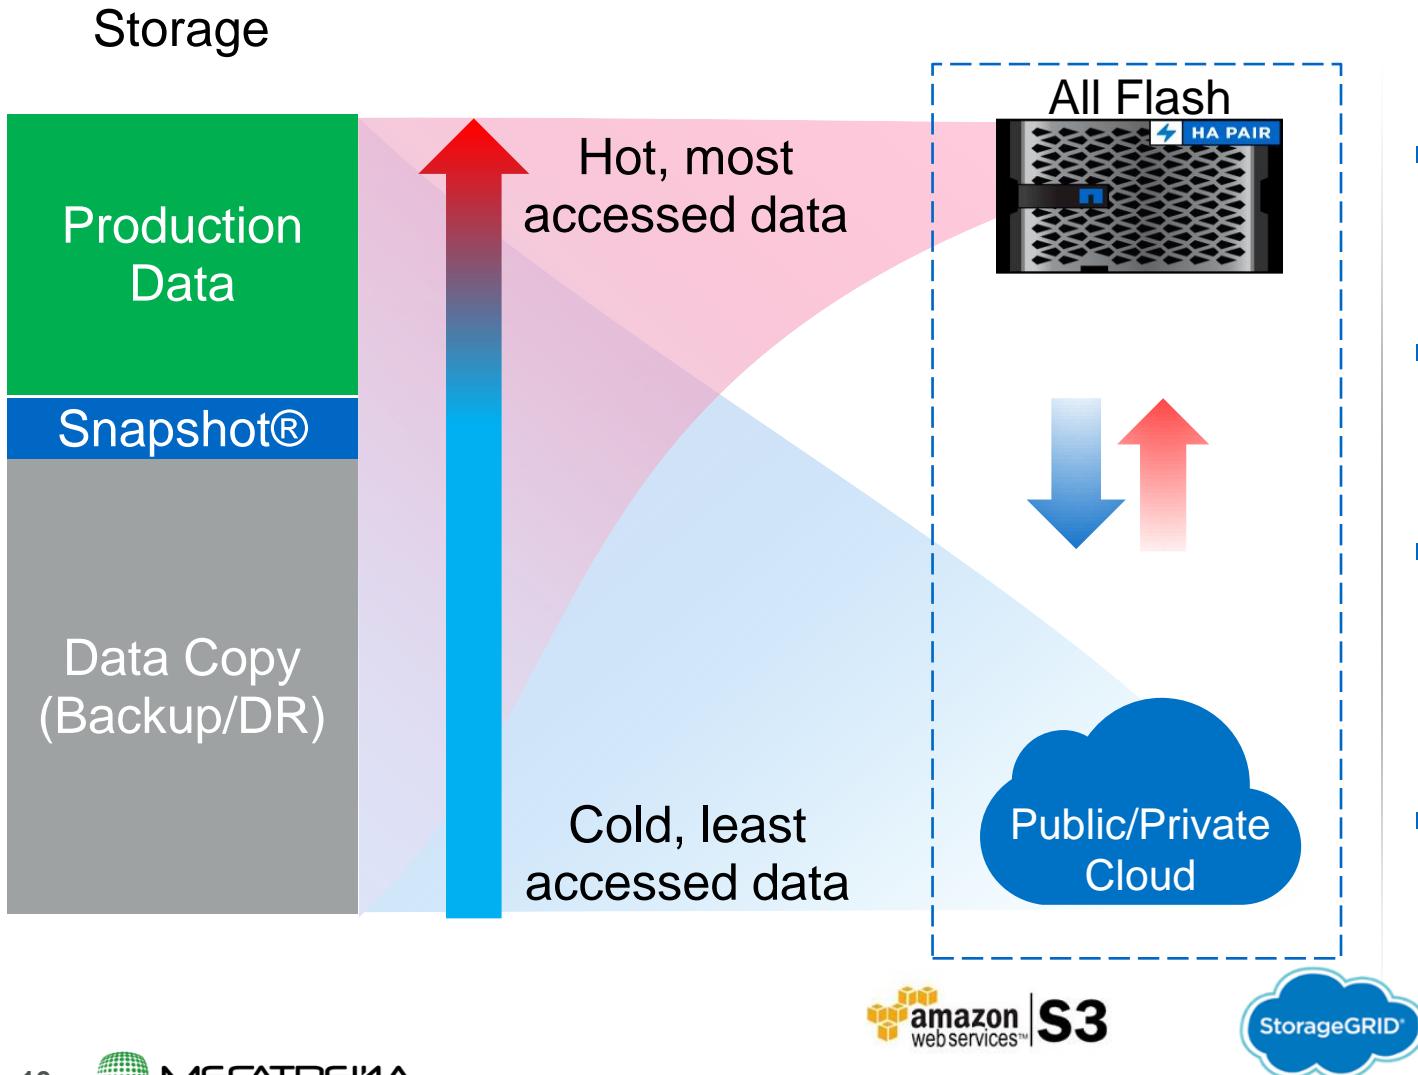
- Идея объединить облако и All Flash массивы в один пул
- Данные делятся на:
 - Hot data - остается на all-flash массиве
 - Cold data - перемещается в публичное облако (AWS) или приватное облако (StorageGRID Web Scale)
- Сохраняет производительность all-flash при одновременном снижении общей стоимости хранения
 - До 40% экономии на ТСО (стоимость эксплуатации)
 - Простота настройки и управления: никаких сложных политик

Расширьте границы вашего ЦОДа



- Требуемые емкости хранилища растут быстрее, чем бюджеты на IT
- Одни данные сегодня востребованы, а завтра нет
- Холодные данные занимают место, но не генерируют доход

Расширьте границы вашего ЦОДа



- Hot data остается на flash; cold data уходит в облако
- Автоматически отслеживается текущее свойство данных
- Данные переезжают без прерывания работы пользователей или приложений
- Данные доступны по запросу

Storage Tiers

[+ Add Aggregate](#)

PERFORMANCE TIER



5.00 GB used of 89.31 GB(5%)

84.31 GB

EXTERNAL CAPACITY TIER



192.75 MB used

AGGREGATES

[✓ aggr123 on the node sti51-vsimg-ucs107f](#)

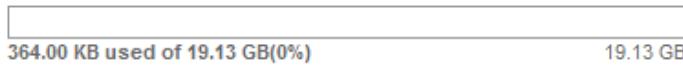
149.46 MB used of 19.13 GB(1%)

18.98 GB

SPACE SAVINGS



3.77 : 1

[✓ sti51_vsimg_ucs107e_aggr1 on the node sti51-vsimg-ucs107e](#)

364.00 KB used of 19.13 GB(0%)

19.13 GB

[✓ sti51_vsimg_ucs107f_aggr1 on the node sti51-vsimg-ucs107f](#)

22.02 MB used of 19.13 GB(0%)

19.11 GB

ADD-ON FEATURES ENABLED

External Capacity Tier



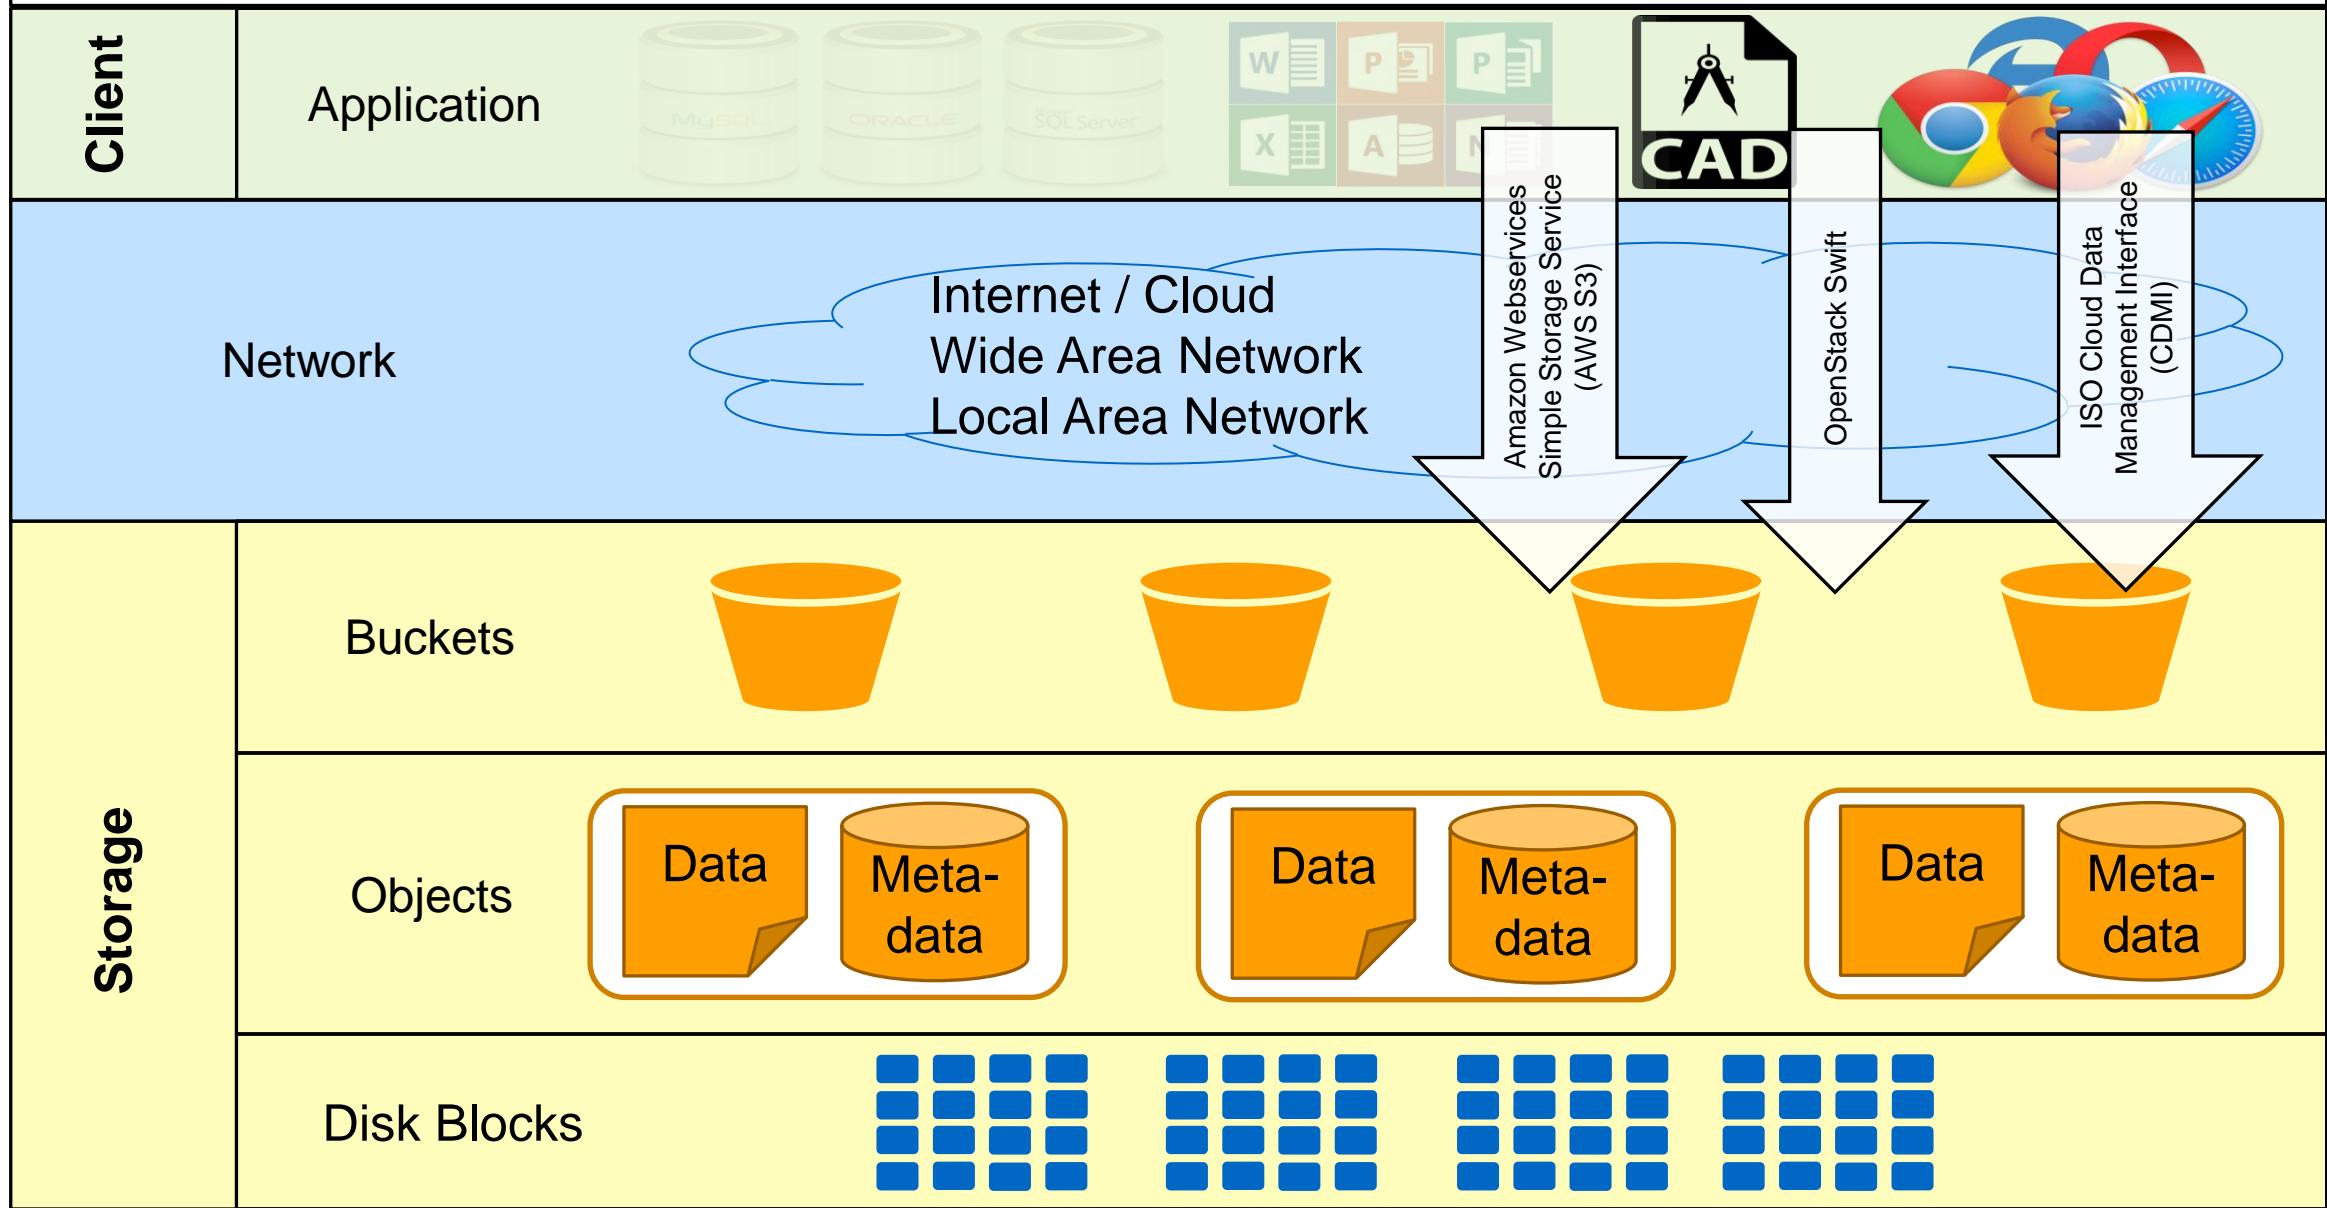
192.75 MB used

External Capacity Tier



0 bytes used

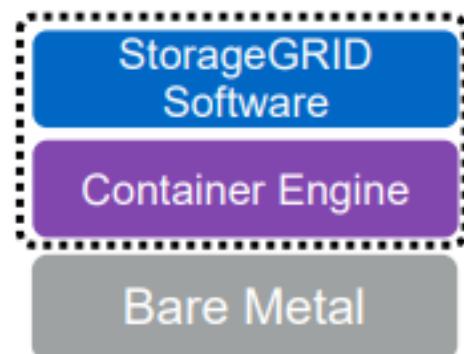
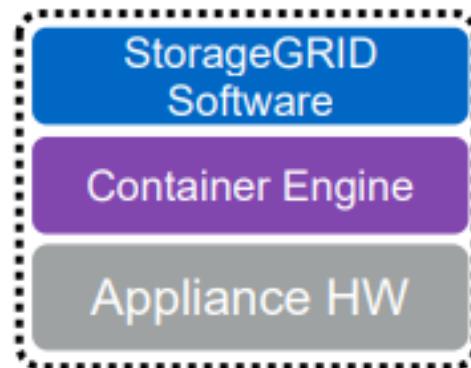
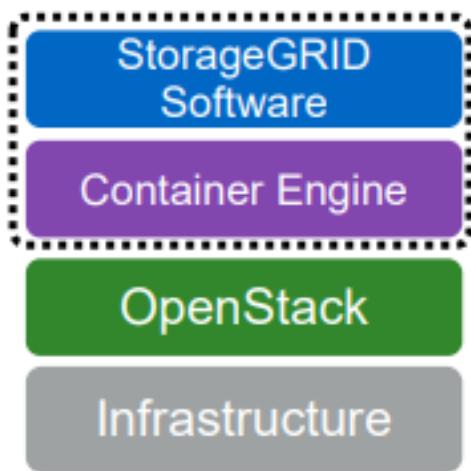
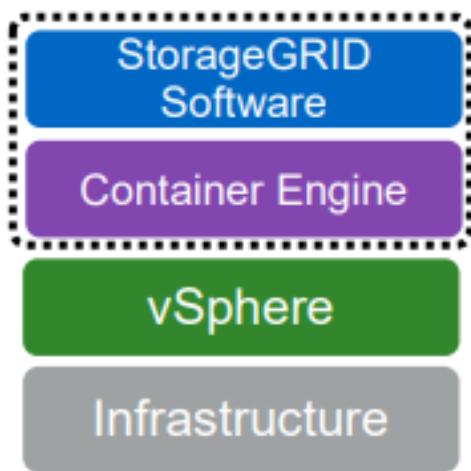
Object Storage



Deployment Options

Mix & Match possible, clear decoupling for HW and SW

NetApp Deliverable



- **VMware-based**
 - ESXi server with external array
 - FAS, E-Series, or 3rd Party
- **OpenStack-based**
 - Consumes Cinder Storage
 - FAS, E-Series, or 3rd Party
- **Appliance-based**
 - NetApp engineered compute, software, networking & storage
 - "Ready to go"
- **Docker-based**
 - Deploy on RHEL/CentOS servers in Docker Container
 - Internal Disks or external array
 - BYOH appliance

StorageGRID Webscale appliance



SG5612



SG5660

- Delivers a Storage Node in an HW appliance
 - Appliance with 4TB, 6TB, 8TB or 10TB drives
 - [SG5612](#), 2U / 12 drives (48/72/96/120TB Raw capacity)
 - [SG5660](#), 4U / 60 drives (240/360/480/600TB Raw capacity)
 - Scale out by adding appliances
- Simplifies deployment and expansion
 - Preconfigured and ready for grid expansion
 - Virtual and appliance nodes interoperate in same grid
 - #TB capacity extention by adding Storage Nodes
 - 500M objects per Storage Node
 - Enhanced throughput and Ingest performance by adding Storage Nodes

Как работает FabricPool?

Как это работает

How FabricPool works

- Подключаем облачные хранилища (AWS, SGWS) к агрегатам all-flash
- Настраиваем политики тиりнга на вольюмах
- Политики тиринга:
 - **Snapshot-only** (стоит по умолчанию)
Cold Snapshot® блоки, не используемые активной файловой системой, переезжают на холодный уровень
 - **Backup**
Данные перемещаются напрямую в capacity tier
 - **Off**
Data is not moved to the capacity tier

Storage Tiers

[+ Add Aggregate](#)

PERFORMANCE TIER



5.00 GB used of 89.31 GB(5%)

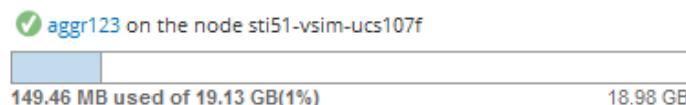
84.31 GB

EXTERNAL CAPACITY TIER



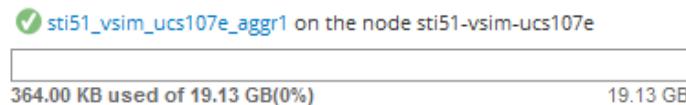
192.75 MB used

AGGREGATES



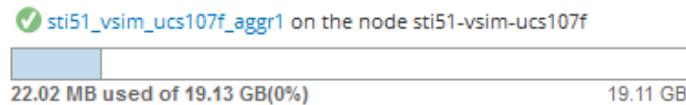
149.46 MB used of 19.13 GB(1%)

18.98 GB



364.00 KB used of 19.13 GB(0%)

19.13 GB



22.02 MB used of 19.13 GB(0%)

19.11 GB

SPACE SAVINGS



ADD-ON FEATURES ENABLED

External Capacity Tier



192.75 MB used

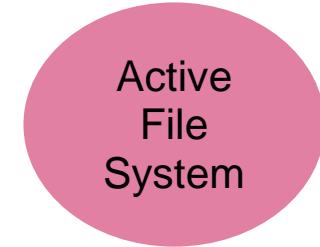
External Capacity Tier



0 bytes used

Tiering data: performance to capacity tier

How FabricPool works



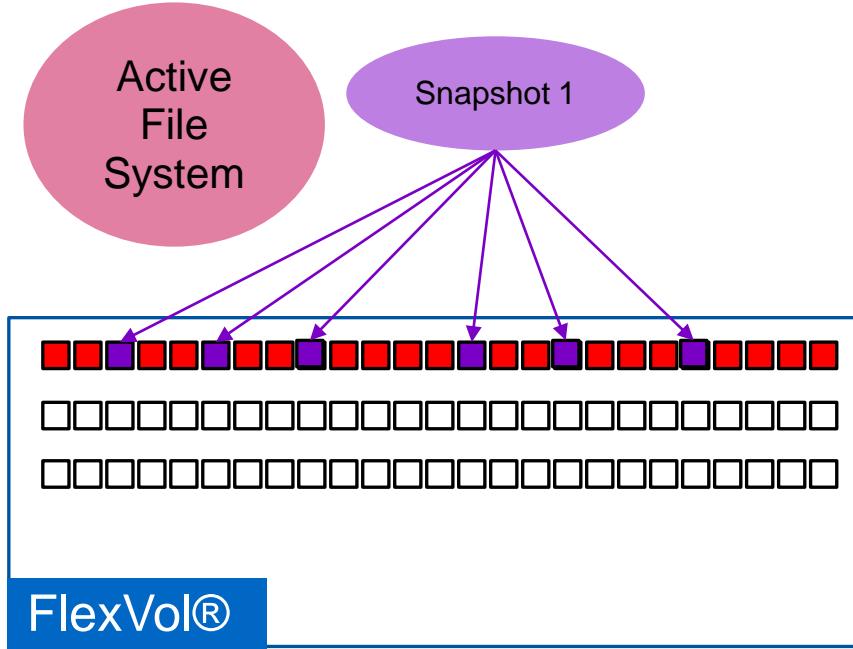
1

Data blocks are written with **temperature value**: hot (■)



Tiering data: performance to capacity tier

How FabricPool works

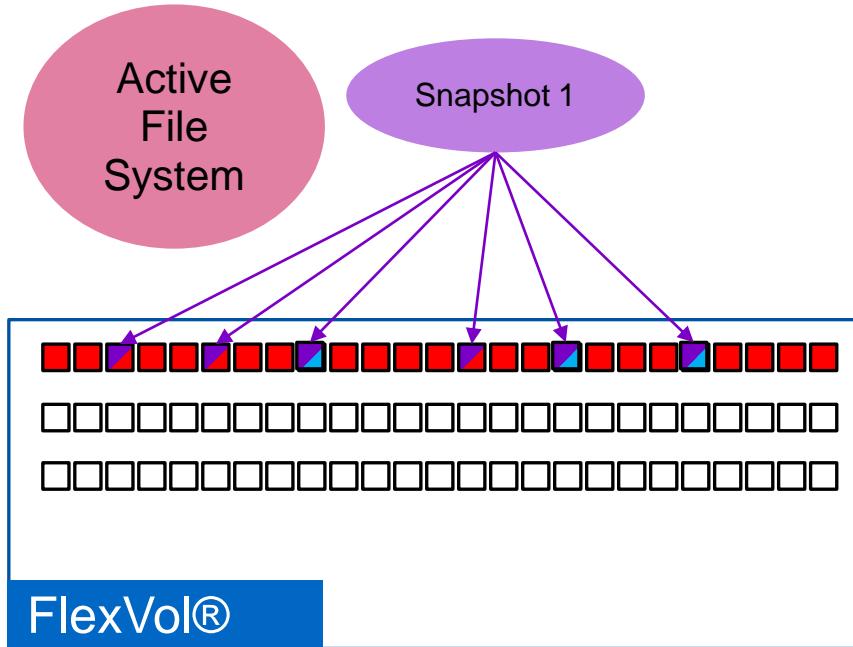


- 1 Data blocks are written with **temperature value**: hot (■)
- 2 Snapshot® copy is initiated
- 3 Overwritten blocks (■) in the active file system become locked to the Snapshot copy



Tiering data: performance to capacity tier

How FabricPool works

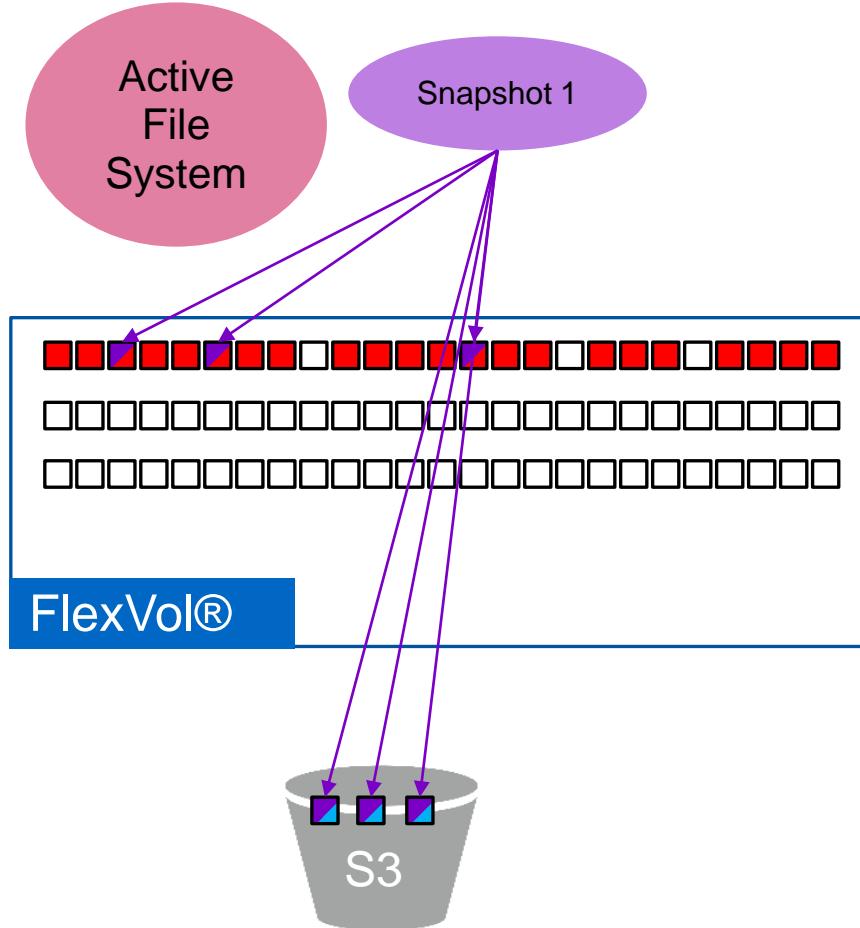


- 1 Data blocks are written with **temperature value**: hot (■)
- 2 Snapshot® copy is initiated
- 3 Overwritten blocks (■) in the active file system become locked to the Snapshot copy
- 4 **Temperature scan** monitors the activity of each block and decreases the temperature value during every scan



Tiering data: performance to capacity tier

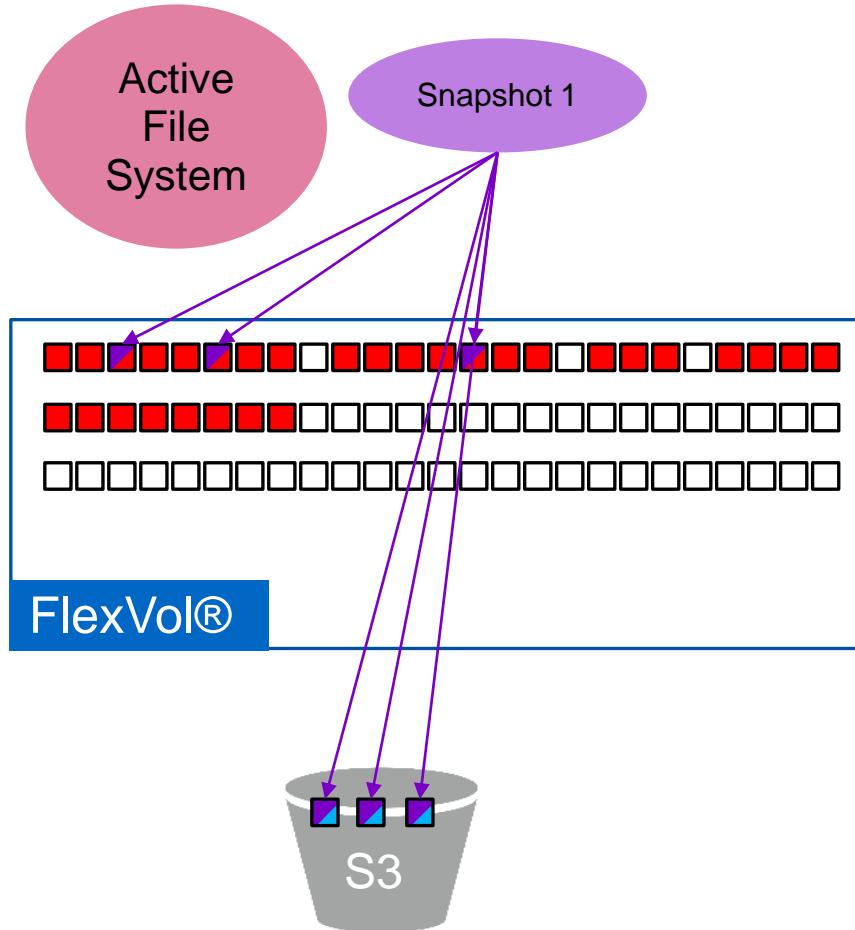
How FabricPool works



- 1 Data blocks are written with **temperature value**: hot (■)
- 2 Snapshot® copy is initiated
- 3 Overwritten blocks (■) in the active file system become locked to the Snapshot copy
- 4 Temperature scan monitors the activity of each block and decreases the temperature value during every scan
- 5 Tiering scan collects cold blocks (■), and packages them into 4MB objects and moves the objects to the capacity tier
 - 1,000 4KB blocks (4MB object)
 - Aggregate used % > 50%

Tiering data: capacity to performance tier

How FabricPool works

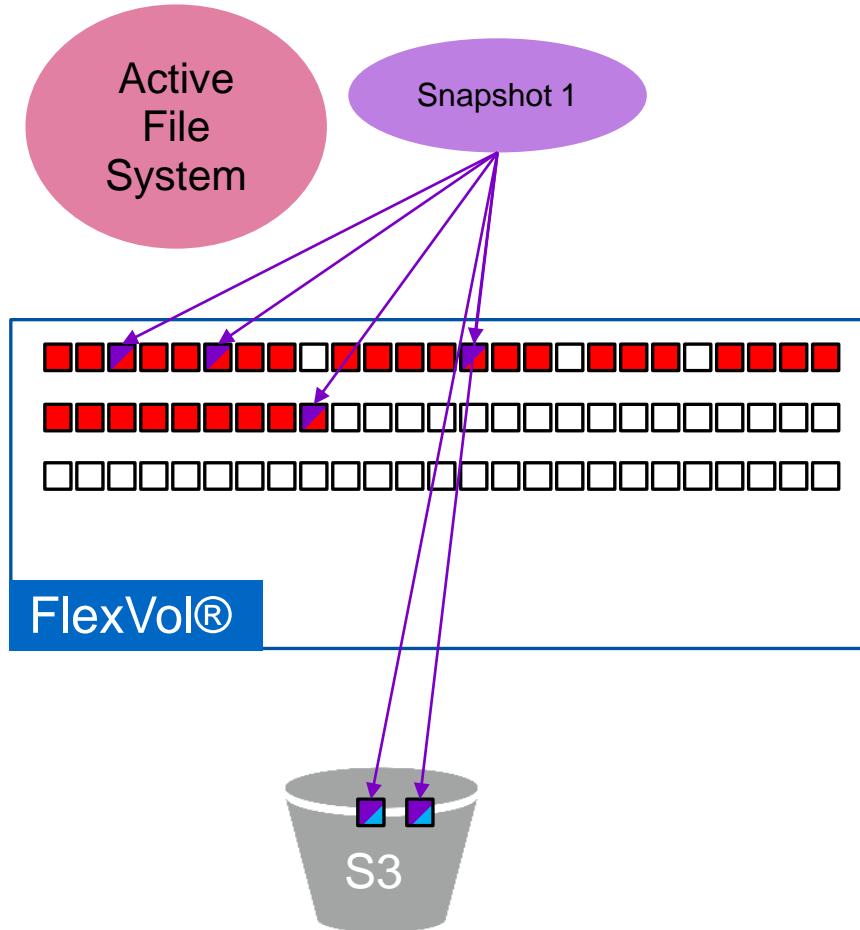


1

Data block **in capacity tier** is read (clone, restore, and so on)

Tiering data: capacity to performance tier

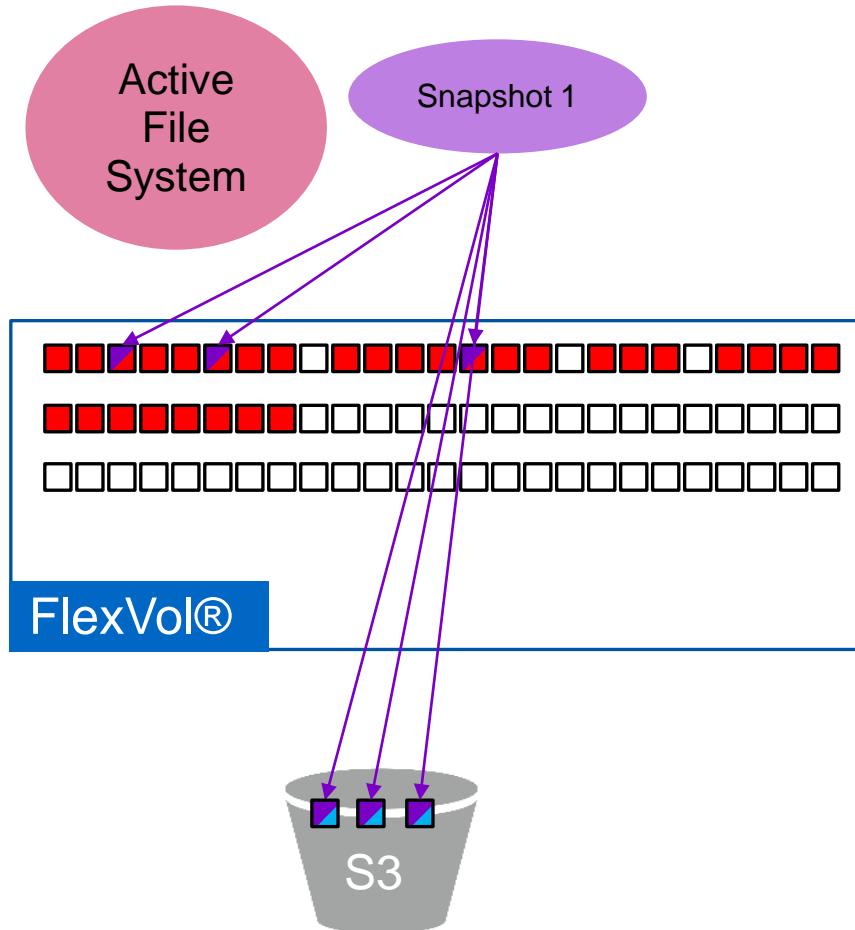
How FabricPool works



- 1 Data block **in capacity tier** is read (clone, restore, and so on)
- 2 Cold block () is fetched using a **get** operation and **moved to the performance tier** and made hot ()
 - Get operation can also fetch required 4KB blocks and not necessarily the entire 4MB object
 - (4MB object is deleted only when all the 4KB blocks are moved back to performance tier)

Tiering data: capacity to performance tier

How FabricPool works



- 1 Data block **in capacity tier** is read (clone, restore, and so on)
- 2 Cold block () is fetched using a **get** operation and **moved to the performance tier** and made hot ()
Get operation can also fetch required 4KB blocks and not necessarily the entire 4MB object
(4MB object is deleted only when all the 4KB blocks are moved back to performance tier)
- 3 Performance tier blocks are temperature scanned and are **moved to capacity tier** if they become cold

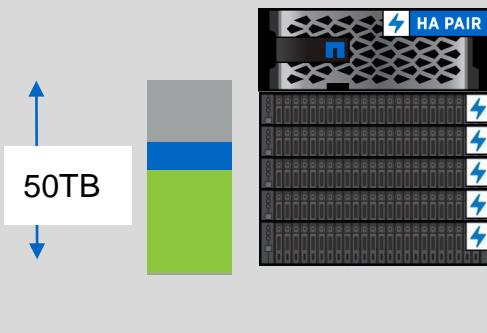
Варианты применения

Освобождаем место на основной СХД

Snapshot only

Before

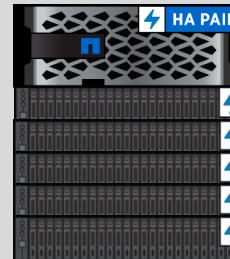
Primary Cluster



After

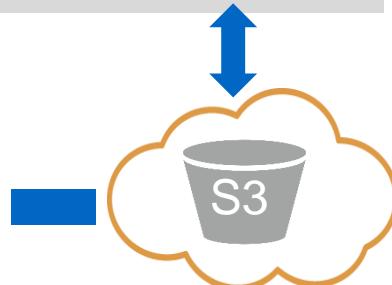
Primary Cluster

50TB



- Provisioned Storage
- Used Storage
- Snapshot® data

5TB

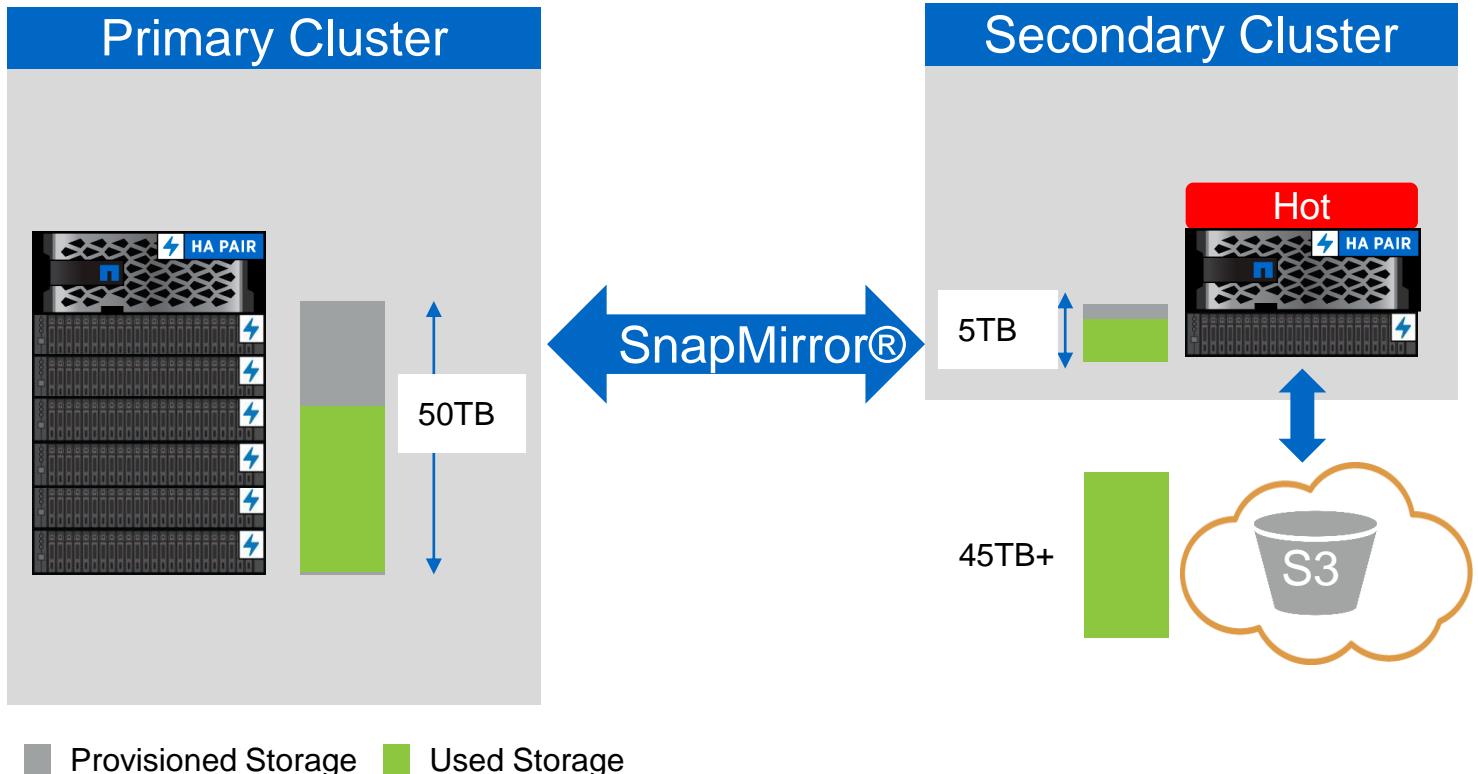


Benefits

- Default volume tiering policy
- Optimize flash usage
- Increase Snapshot copies (maximum remains 255 Snapshot copies)

Уменьшаем размер дополнительной СХД

Backup

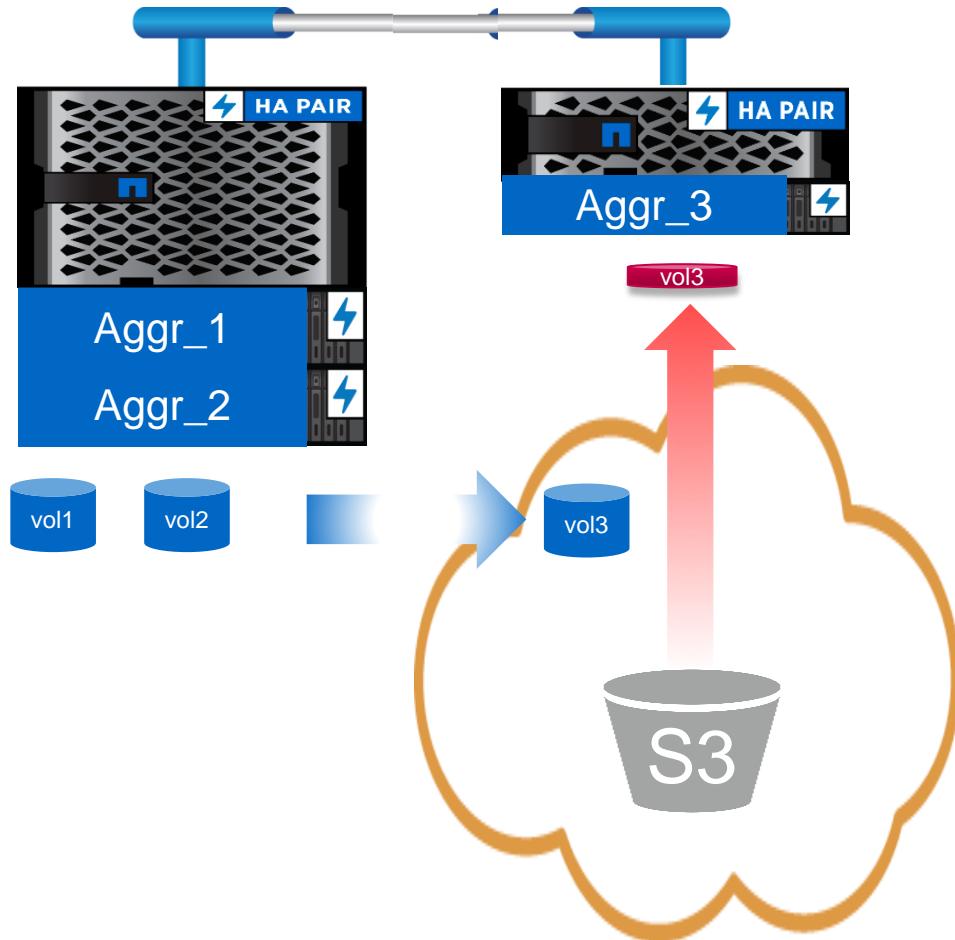


Benefits

- Expand capacity on secondary (data protection) cluster
- Reduce footprint on secondary cluster
- Existing data protection policies work seamlessly

Backup на уровне вольюмов

backup plus volume move



Benefits

- Data moves to capacity tier
- ONTAP® system metadata stays on performance tier
- When read, capacity tier data becomes hot and is moved to the performance tier

Поддерживаемые платформы и лицензирование



All Flash FAS

A700S, A700, A300, A200
AFF8080, AFF8060, AFF8040



FAS*

FAS9000, FAS8200, FAS2650, FAS2620
FAS8080, FAS8060, FAS8040

*Только SSD агрегаты могут использовать FabricPool



ONTAP® Cloud

Performance Tier: gp2 and st1 volumes
Capacity Tier: AWS S3



Object storage

AWS S3, StorageGRID® Webscale 10.3+



Спасибо!