

NetApp Cloud Tiering and Inactive Data Reporting



Executive Summary



The future of the enterprise was always going to be digital. However, no one anticipated how rapid this future would be, and the recent worldwide work disruptions have only accelerated this transformation. These circumstances have significantly increased the reliance organizations have on digital technologies, and as a result, data footprints and costs have grown larger than anyone could have anticipated. This brings about a number of challenges when it comes to the long-term storage of such data, especially since such data is considered “cold” due to no longer being regularly accessed.

The ability to identify this cold, inactive data and store it in a low-cost object storage platform is a key priority for enterprise organizations to prevent total cost of ownership (TCO) from spiraling out of control. This guidebook takes a look at the accelerating data growth, the different categories of data, and how NetApp Cloud Tiering and Inactive Data Reporting make it possible to easily identify cold data and define appropriate data storage strategies so that corporate IT budgets are not exasperated by crippling storage costs of inactive data, over their entire lifecycle.

Table of Contents


Introduction: The New Shape of Data	4
What Is Cold Data?	6
How Can Cloud Tiering Service Help?	8
Assessing Cold Data Capacity	9
Key Use Cases and TCO Estimation	10
Recommended Approach	11
Conclusion	12

Introduction: The New Shape of Data

Over the last decade, the digital transformation has slowly been driving the next wave of the digital revolution. Today, every organization uses some kind of a digital technology platform to power the nucleus of their business operations. Often, the innovative use of next-generation digital technology is a fundamental must-have for an organization to continue to grow and be successful, in an increasingly digital world for both consumers as well as organizations.

This concept of digital transformation revolves around integration of digital data into everything we do. In effect, users (and devices) generate front end data which is digitally monitored and tracked at the “Edge” for immediate analysis and decision making. A summary of this data is then sent off to the “Core,” often powered by a cloud computing platform (public or private) for further analysis so that the organizations can derive meaningful intelligence.

As a part of this analysis, the whole process creates more data that will then be used by various digital algorithms to predict the future demands and needs so organizations can better serve their target consumers.






Take the concept of a connected car as an example. Numerous IoT sensors installed in the car will generate thousands of gigabytes of data an hour, covering the telemetry, the health and performance of various components of the car as well as its surroundings as observed by the cameras automatically. This data will then be aggregated at the edge: A compute and storage subsystem onboard the car which analyses these data and important metadata (data about data) is then passed over to the Core: i.e., a traditional data center or a cloud data center, for further analysis. The car manufacturer / fleet management organization can then use this intelligence (data) collected over-the-air from all of its fleet to provide tailored over-the-air updates and data management services to each of its customers such as performance improvements, on-demand system updates, or integration with other third-party platforms, such as an insurance service for example, some of which can act as additional revenue streams. Throughout the entire process, vast sums of data are created and processed both at the Edge, as well as at the Core.

As the ways and means of this digital transformation evolve due to new technologies such as IoT, AI, and the yet-to-be-invented next big thing, one thing is guaranteed for sure: there will be more and more digital data generated in the future. But while some of this data will be transient in nature, the vast majority of it will be permanent. If you think back on the connected car example above, the data collected needs to be stored long term, for purposes such as compliance, analysis, or monetization.

IDC estimated that [more than 59 ZB of data would be created, captured, copied and consumed during 2020 alone](#), aided by the abrupt increase due to the COVID-19 pandemic driving increased consumption of digital technologies. This is in-line with a similar [Cisco analysis](#) that predicted a massive buildup (4.5ZB) of world's IP traffic by 2022 due to the increased use of IoT traffic, video and new users, and with a [joint Seagate & IDC study](#) that predicts the world's collective data storage requirements to grow from 45ZB back in 2019 to 175ZB by 2025.



How can you save on the costs for storing this growing amount of data?

This exponential data growth brings a number of challenges to many organizations. Most organizations struggle to keep up with this level of storage needs due to the cost of the underlying storage and management infrastructure requirements. Enterprise storage solutions are typically expensive, given the high-performance requirements they have in order to serve the needs of regularly accessed data. Therefore, using such solutions to store all data can be prohibitively costly for many organizations, especially when as much as 70% of the data on hand can be considered “cold,” i.e., not accessed regularly.



What Is Cold Data?

All digital enterprise data tends to traverse through the typical data life cycle stages (see sidebar). As it traverses through these stages, the characterization of data changes in terms of its relevance and value. Freshly created data is generally more important and is of higher value due to the up-to-date-ness of it and often this data is classified as “hot” or “warm” data. Hot and warm data tend to be accessed frequently by various applications and enterprise users for frequent processing purposes.

Data Lifecycle Management Basics

Data Lifecycle Management refers to the process of understanding the various stages that data goes through during its existence. Key phases of a typical data lifecycle include:



Stage 1: Data generation

Creation of data through acquisition of existing data, manual entry of new data, and capture of data generated by various systems.



Stage 2: Data processing

Processing of data created to reduce the noise and discard irrelevant data. During this stage, data is typically accessed frequently and requires to be stored locally in places like the edge or at the core, such as an enterprise data center.



Stage 3: Data storage and consumption

Active storage of processed data for an organization’s objectives and operations. Similar to the previous stage, data within this stage is typically stored at the core in highly performant storage (data center or cloud-based repository).



Stage 4: Data archival

The active use of data has completed, and the data is typically stored for long term retention and storage efficiency reasons. Archived data is typically stored in low-cost storage tiers at the core, such as tapes, or in the cloud within cloud-based object storage tiers, just in case it’s required in the future.



Stage 5: Data purging / retirement

In this phase, data which is no longer needed to be maintained is permanently deleted.




Once such hot and warm data has been processed and analyzed, and has served its initial purpose of providing meaningful intelligence, that data is often required to be stored for a lengthy period of time, in an infrequently accessed cold state and there are a number of reasons for this. Many organizations, such as financial institutions and public sector organizations, have various regulatory and compliance requirements forcing them to retain data from six months up to 25+ years. In some cases, data may need to be stored indefinitely too.

While such requirements can vary across regions around the world, one good example is the clearly defined data retention requirements from [the UK's Financial Conduct Authority \(FCA\)](#), which show the data retention requirements mandatory for every local and international financial organization that operate within the UK. There are also [similar requirements from the USA's federal reserve](#).

In addition to such regulatory requirements, many organizations will store and retain such cold data for lengthy periods of time due to various operational requirements. Historical data, such as the data from a completed project, or historical recordings of a share price movements can be used effectively for data mining purposes in the future and technologies such as big data and AI can work hand in hand with such historical data to predict future behaviors and business opportunities. That makes such cold data valuable for every organization and retaining it in an electronic

media that can be accessible on-demand is a must have requirement. Another type of cold data with long retention requirements is data kept for security, backup, and disaster recovery (DR) purposes. Such data is often stored at an alternative location from their main copy, and this data typically belongs to the “archival stage” of the enterprise data life cycle. In an increasingly digital world, digital data is the biggest asset that any organization has. Keeping proper backups of such data in a secure, offsite location is an existential requirement for many in the time of a crisis. Such crises can include accidental deletion or major security incidents such as ransomware attacks that render the primary hot / warm dataset inaccessible.

Cold data such as backup or DR copies as well as historical data retained for future use can constitute up to 70% of a typical enterprise organization's storage requirements. Such data types don't require regular, frequent access by nature, and that makes them suited to be stored on high-capacity, low-cost object storage solutions such as Amazon S3, Azure Blob, and Google Cloud Storage. This can significantly reduce the total cost of ownership of such data throughout their entire lifecycle.



How Can Cloud Tiering Service Help?






NetApp Cloud Tiering is a Software as a Service (SaaS) offering that helps on-premises ONTAP systems to benefit from the low-cost object storage solutions in the cloud by leveraging NetApp's FabricPool technology. Cloud Tiering extends high performance flash tiers located on premises to the cloud by seamlessly moving cold data to high-capacity, durable, low-cost object storage tiers, with no impact to the front-end applications and users of that data.

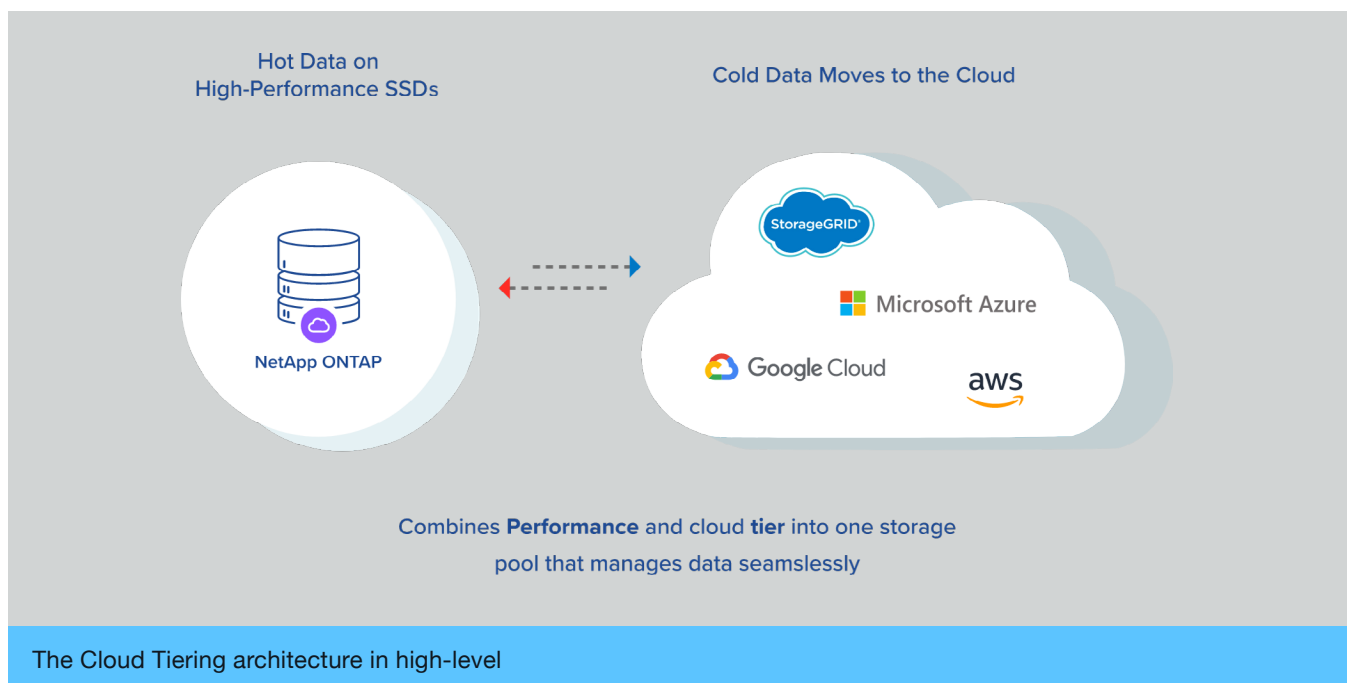
With Cloud Tiering, active data (both hot and warm data) remains on the high-performance tiers in the data center to meet the performance needs of the application, while cold, inactive data is automatically identified and tiered off to an object storage platform, freeing up valuable capacity on the on-premises storage array.

ONTAP storage efficiencies such as deduplication and compression are also maintained within the cloud tier which ensures the cost savings are optimized on the object storage tier due to minimum storage consumption.

Cloud Tiering can help enterprise customers leverage their existing cloud storage capacity, often included as a part of existing [Microsoft Enterprise Agreements](#) or [AWS Enterprise Discount Program](#) (or the like) to significantly reduce their TCO of the on-premises storage infrastructure.

Tiering to the cloud in general helps customers address number of typical challenges such as:

-  Allocating CAPEX budget for data center expansion.
-  Meeting cloud strategy goals in a timely manner.
-  Complexity in migrating workloads to cloud due to impact on application architecture and business workflows.
-  Identifying ideal workloads for cloud migration.
-  Unpredictable data growth and difficulty in forecasting storage requirements.
-  Reducing storage total cost of ownership (TCO).



Assessing Cold Data Capacity

For any enterprise, being able to clearly identify the “cold” portion of their data can provide a clear advantage when it comes to maintaining the TCO of the data retention. However, this is also one of the hardest things to do for many; from the outset, it is not that easy for administrators to clearly identify and segregate hot and warm data from cold, regularly inactive data.

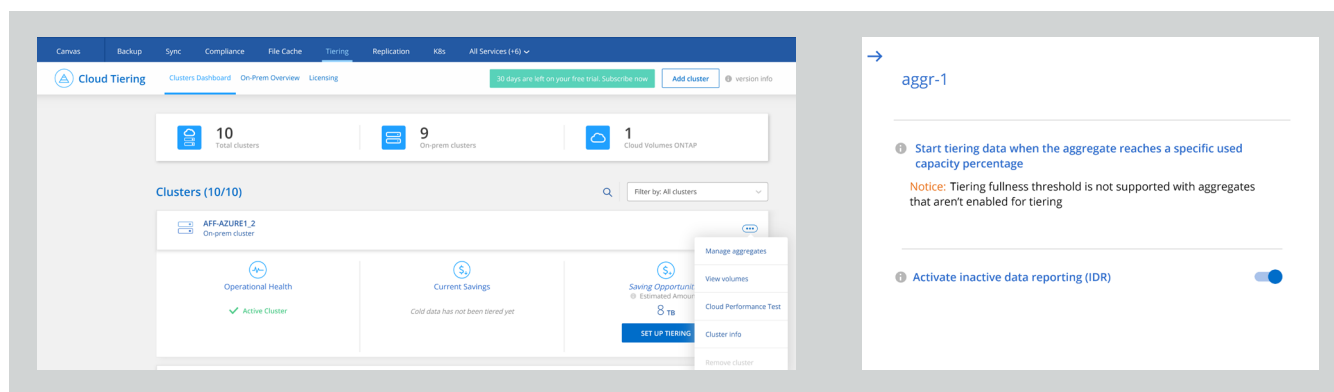
Luckily, all NetApp customers who are using ONTAP data management solutions can consider this a non-issue due to the built in Inactive Data Reporting (IDR) feature. Available since ONTAP version 9.4, IDR provides an easier way to determine how much inactive, cold data is available on each ONTAP aggregate and volume. This data can then be tiered off seamlessly to a cheaper cloud-based object storage solution using NetApp Cloud Tiering, saving significant costs.

IDR is automatically enabled on all SSD aggregates on ONTAP 9.6 (and above), as well as on all FabricPool aggregates from ONTAP 9.4 onwards. Starting with ONTAP 9.6, IDR can be enabled on all other aggregates (including HDD aggregates) as well.


Once enabled, IDR monitors the underlying data access across a fixed 31-day cooling period to determine which data is considered inactive. Starting with ONTAP 9.8 the cooling period used by IDR can be adjusted. Once identified, the size of this data can be visible on a per-aggregate and a per-volume basis.

Using IDR with Cloud Tiering

IDR can be enabled from within the NetApp Cloud Tiering service, integrated into [NetApp Cloud Manager](#) (Tiering tab -> Clusters Dashboard -> Manage aggregates) as shown below.



The Cloud Tiering service also provides visibility into IDR findings, displaying the amount of inactive cold data at both aggregate (manage aggregates page) and volume level (tier volumes page) as shown below.



For additional information on Cloud Tiering please refer to the [NetApp Cloud Tiering documentation](#).

Being an ONTAP feature, IDR can also be fully managed using the below CLI commands:

- To enable IDR on an aggregate:

```
storage aggregate modify -aggregate  
aggr1 -is-inactive-data-reporting-  
enabled true
```
- To show inactive data for an aggregate:

```
storage aggregate show-space -fields  
performance-tier-inactive-user-data,  
performance-tier-inactive-user-data-  
percent
```
- To show inactive data for a volume:

```
volume show -fields performance-tier-  
inactive-user-data, performance-tier-  
inactive-user-data-percent
```

It should be noted that IDR reports all cold data found (both snapshot and user data) and once cold data tiering is enabled for a volume, the data marked as inactive is moved to the cloud tier based on the tiering policy associated with the volume.

Note that IDR cannot be enabled on non-FabricPool compatible aggregates, such as the root aggregates, FlashPool aggregates, and SnapLock aggregates.

Key Use Cases and TCO Estimation

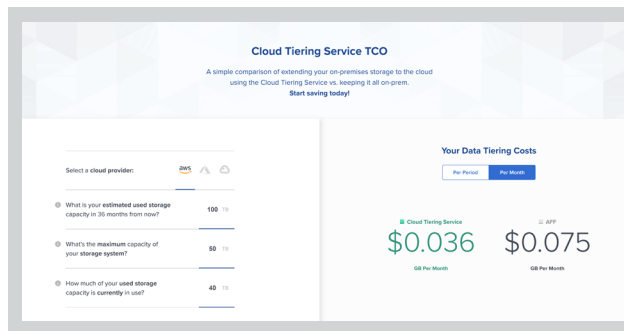
There are [various typical enterprise use cases](#) for Cloud Tiering. Purchasing a new NetApp AFF system can now be significantly cheaper due to reduced hardware footprint thanks to Cloud Tiering. The same applies for expanding existing AFF storage systems either due to running out of space, or due to net new capacity requirements stemming from new business initiatives. Thanks to Cloud Tiering, capacity on existing storage arrays can be reclaimed and the need for additional SSD shelves and the associated CAPEX commitment can significantly be reduced, cutting ongoing data center operational costs further.

In addition, NetApp Cloud Tiering makes it a breeze to embrace a cloud strategy. A top priority for many CTOs in the enterprises today in order to reduce overall IT costs. Cloud Tiering helps to achieve a hybrid cloud strategy and cost savings without the usual caveats such as the application refactoring costs and data migration costs. Refer to [this example of how an Ivy league university used Cloud Tiering](#) to clear over 100TB of on-premises data while saving money on hardware, electricity, and space.

Perhaps one of the biggest use cases for Cloud Tiering comes from backup and DR. Cloud Tiering enables organizations to significantly reduce their storage costs when it comes to storing an identical copy of their production data in a DR location. The storage hardware footprint required in the DR site can now be significantly reduced by tiering off those DR copies into cheaper object storage using Cloud Tiering. Likewise, snapshot and backup data can also be easily tiered off, significantly reducing the need for retaining long term backup data on premises on expensive storage.

[Read how Festo, an industrial giant in Germany, managed to leverage Cloud Tiering to tier a large proportion of their Snapshot data to object storage on the cloud.](#)

One thing that many customers struggle with at the beginning of many cloud projects is to accurately establish the TCO savings as a part of the initial business case. In order to address this requirement, NetApp has now made available a [Cloud Tiering TCO calculator](#) where customers can easily calculate the cost savings to be had upfront.



Recommended Approach

With NetApp Cloud Tiering, customers have the full flexibility to adopt Cloud Tiering across all the typical use cases and start realizing all of those benefits from day zero. Harnessing ONTAP's proven tiering technology, customers can rest assured that their valuable data is in safe hands from the get-go and are not subject to proprietary platform locking from the cloud or object storage service providers.

Using the cloud—the public cloud in particular—to support tiering cold data, is directly related to how confident organizations are with the cloud. As such, for customers starting their journey in the cloud, it is recommended to start simple with the easiest use cases such as the one defined by the “Cold Snapshots” (Snapshot-only) tiering policy. This will ensure that only snapshot copies, that have minimum read requirements and are considered non-production data, are first moved to the cloud tier. As customers start to feel more comfortable, and confidence in the public cloud is built, additional use cases and volume tiering policies can be added to tier off further data on to the cloud tier.


Conclusion

NetApp Cloud Tiering can help existing NetApp AFF and FAS customers with All Flash aggregates to optimize their storage and achieve up to 70% storage capacity savings through the intelligent use of cheaper object storage solutions from the public or private cloud. The Inactive Data Reporting feature further enhances these benefits.

Cloud Tiering can prolong the life of their existing data center hardware by increasing capacity by up to x50 and can reduce TCO of underlying storage on average by 30%. All this is made possible without the need for the usual complexity associated with cloud storage adoption such as application refactoring and data migration. With Cloud Tiering, IT departments can now spend these costs savings and the man-hour savings on supporting valuable new business initiatives instead.

**Start saving on
your on-prem
data costs today.
Sign up for a free
trial of Cloud
Tiering to get
started.**

Start now



Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NA-000-1220