SNIA DEVELOPER CONFERENCE
SDC 25
By Developers FOR Developers

Hyatt Regency Santa Clara, CA
September 15-17, 2025

Smarter Cloud Storage: Optimizing Costs with Tiering & Automation

How to Align Storage with Usage and Save Big

Presenters:

Vignesh Shanbhag, *PM Azure Storage*

Shruti Sethi, *Senior PM Azure Storage*

Agenda

- Cloud Storage Cost Challenge
- Storage Tiers & Their Importance

- Tiering Policies
- Automating Tier Changes

- Key Challenges & Blind Spots
- Tier Optimization for AI Workloads
- Monitoring & Visibility of Data-Usage Patterns

- Best Practices for Cost-Optimization
- Key Takeaways

Microsoft

SDC 25

The Cloud Storage Cost Challenge

Data growth is exploding and storing everything in “hot” storage is unsustainable

90% of data is rarely accessed after 90 days

- **Problem:** Lots of cold data sitting in expensive hot tiers by default
- **Consequence:** Wasted spend on storage that isn’t delivering value
- **Task:** Optimize Performance and Cost Savings

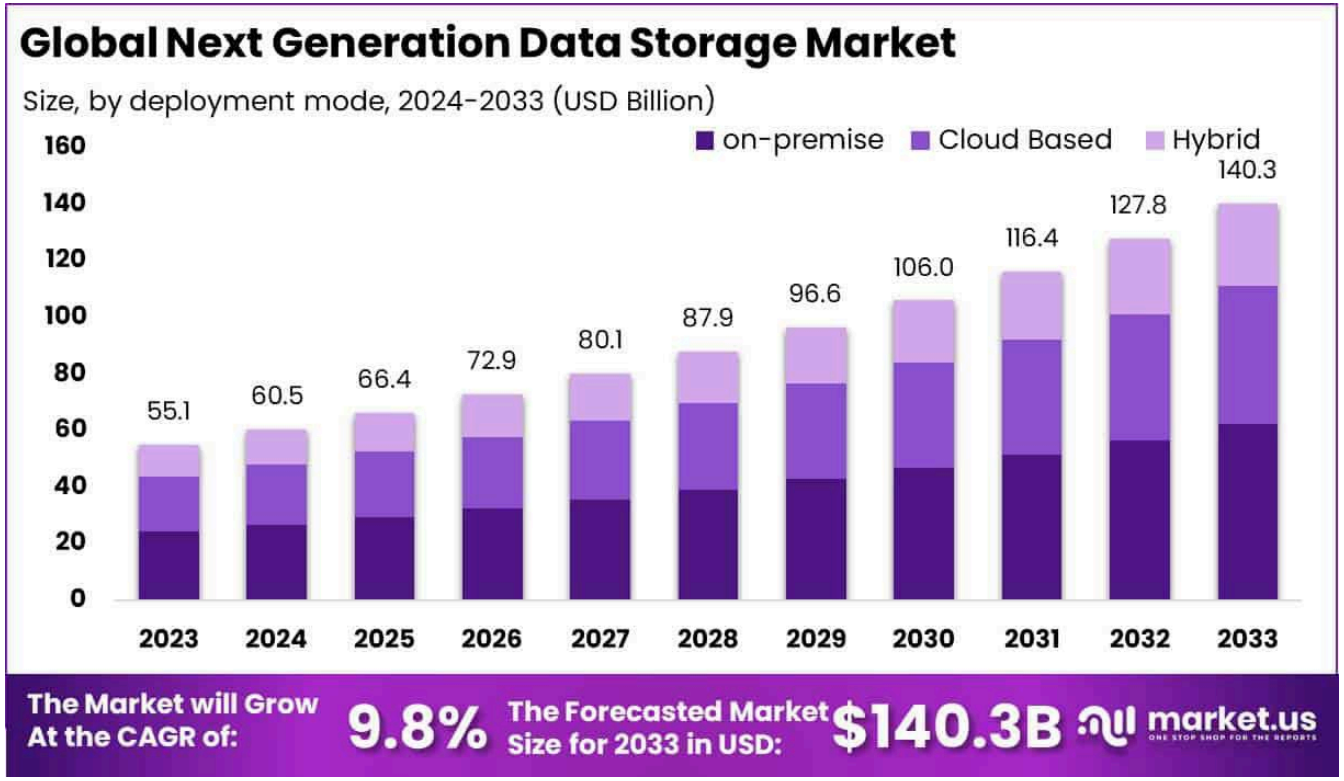

Storage Costs represents 10-30% of total cloud spend for most enterprises

Microsoft

SDC 25

Storage Tiering & Why It Matters





	Premium	Hot	Cool	Cold	Archive
	Low and consistent latency data	Frequently accessed data	Less frequently accessed data	Rarely accessed data - Online	Rarely accessed data – Offline
Per TB Per Month	$150	$21	$15	$3.6	$1
Per 10K Read Operations	$0.0019	$0.005	$0.013	$0.13	$6.5
Per 10K Write Operations	$0.0228	$0.065	$0.13	$0.234	$0.13
Retrieval Times	Milliseconds	Milliseconds	Milliseconds	Milliseconds	Hours

- **Hot vs. Cool vs. Cold vs. Archive:** Warmer the temperature, faster the access, higher $/GB
- Moving rarely-accessed data to lower tiers can cut storage costs by **90%+**
  - Archive tier ~1/10th the cost of hot

Reference

Microsoft

SDC 25

Tiering vs Transactions: Breaking Even

- Objects tiered down to cooler tiers must stay for a minimum duration
- Let's take an example of 1GB of data stored in the Cold tier

Early Deletion Penalty	
Max Early deletion penalty (90 days)	$ 0.0240900
Cold to Hot tier txn (billed as read)	$ 0.0000130
Total Penalty	$ 0.0241030

Transactions in Cold to Break even	
Read transactions	1,854
OR	
Write transactions	462

1GB of data = 562 objects (avg. of 2MB per object)
If transactions on each object is less than 5 reads or 2 writes per object in a 30-day period, it is cheaper to leave it in Cold tier.

Microsoft

SDC 25

Using Tier Move Policies

- Policies define rules for Tier-Down or Tier-Up
- One-Time effort to set policies / rules for tiering
- Policies can be defined by user or via automation
- Can be applied to all existing and new data/accounts/objects

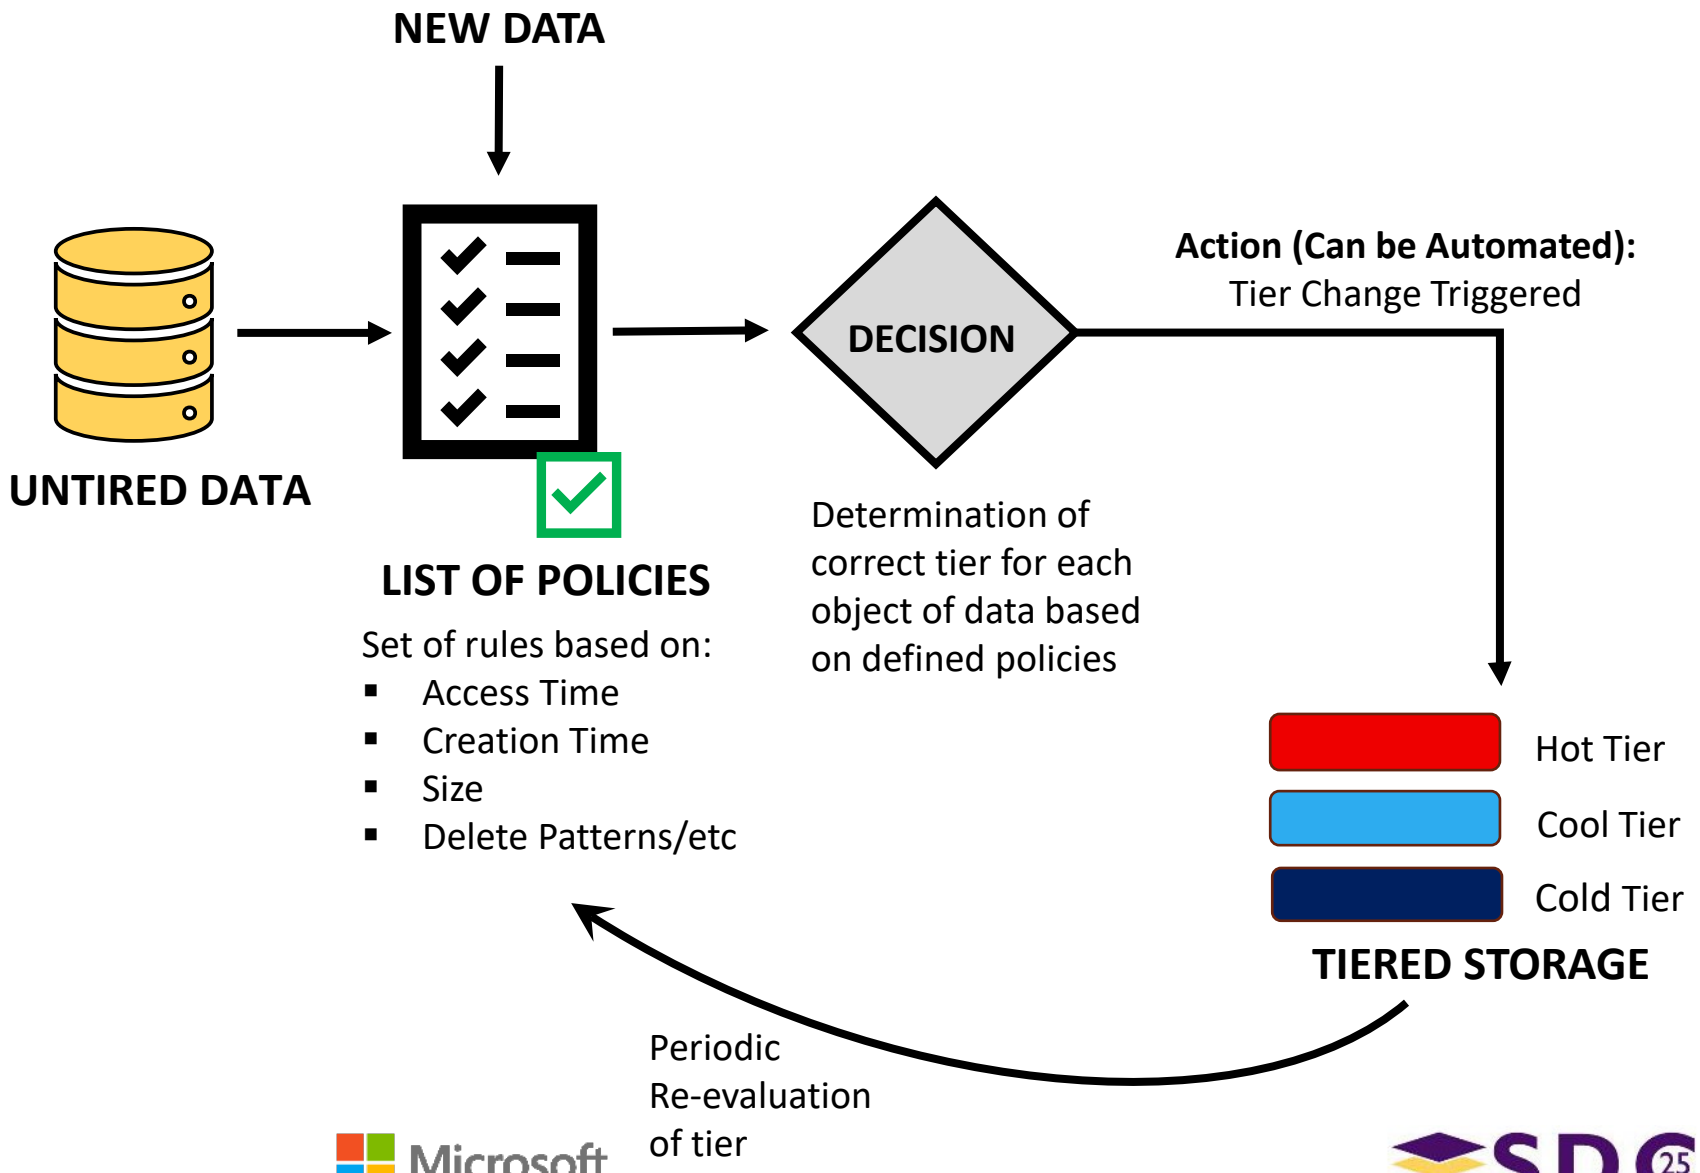

Automating Tier Changes (via Policies)

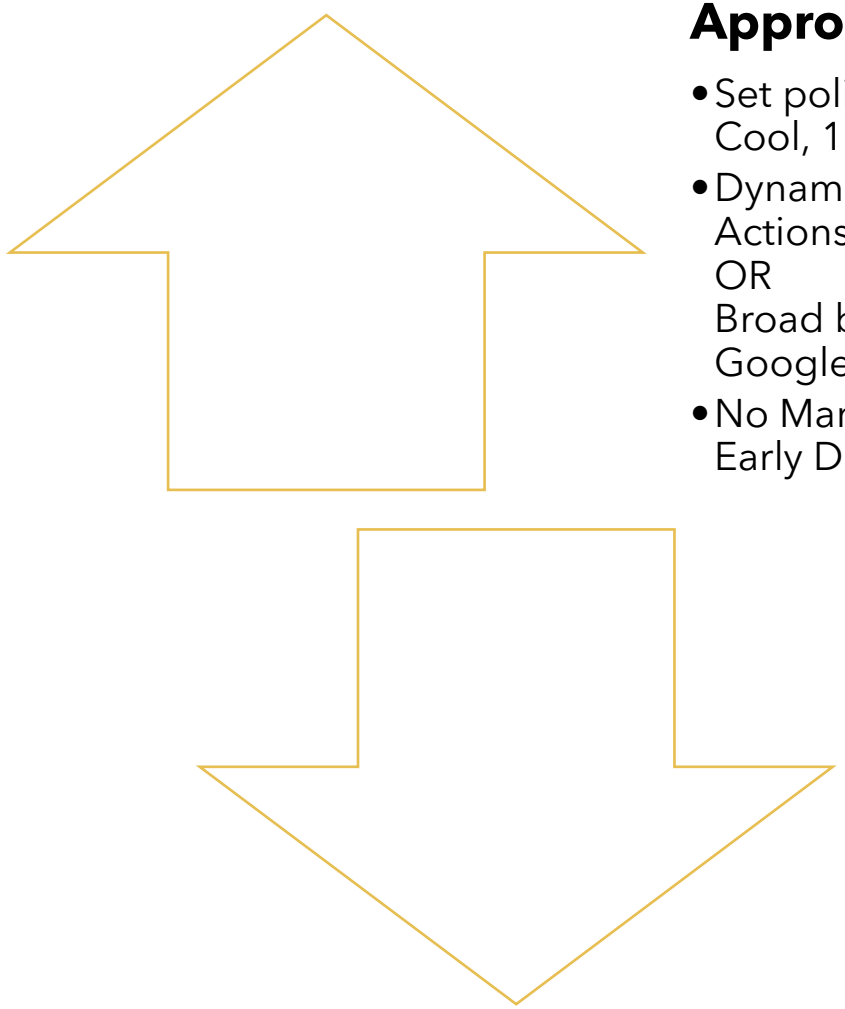

Approach 1: Data Management Policies

- Set policies to Auto-move data as it ages (e.g. 30 days -> Cool, 180 days -> Archive)
- Dynamic conditions and targeting with Azure Storage Actions
  OR
  Broad brush stroke policies with Blob Lifecycle Management, Google Object Lifecycle Policies
- No Management fee. Tier Move transactions are charged. Early Deletion penalties apply

Approach 2 : Intelligent Tiering

- Use Intelligent Tiering services that auto-shift data between tiers based on access patterns
- No need to predefine policies by user – service monitors usage
- No Early Delete Penalties or charges for the Tier move. Management fee for the entire account.
- eg. AWS S3 Intelligent-Tiering, Google AutoClass

Microsoft

SDC 25

Storage Actions in an Example workflow

Challenge: Autonomous vehicle development generates massive volumes of data—LiDAR, camera feeds, radar, GPS, and simulation outputs—requiring efficient storage, organization, and data management.

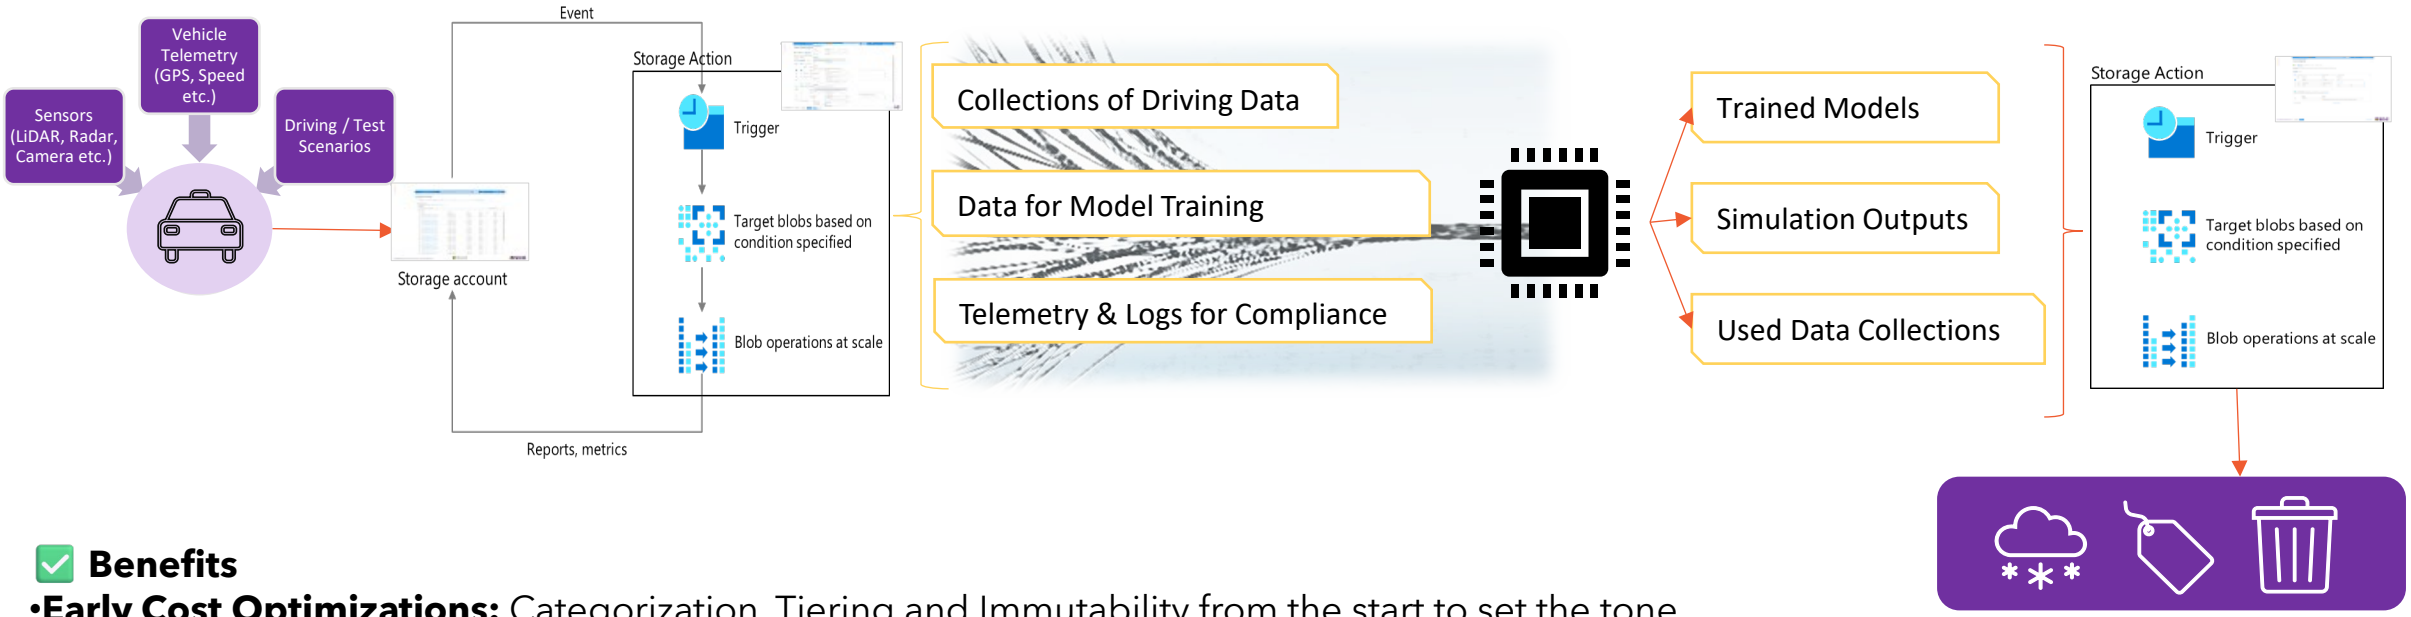

✅ **Benefits**

- **Early Cost Optimizations:** Categorization, Tiering and Immutability from the start to set the tone.
- **No-Code Automation**: Simplifies complex workflows without scripting.
- **Scalable Across Billions of Objects**: Ideal for enterprise-scale data operations.
- **Cost Savings**: Reduces storage costs and operational overhead.
- **Consistency & Governance**: Ensures policy enforcement across regions and teams.

Microsoft

SDC 25

Key Challenges and Blind spots

Access Pattern Related Blind Spots

📈 Lack of Visibility & Monitoring

- Tiering without clear insight into data usage

🧱 Considering Static Data Characteristics

- Treating data characteristics as static over lifetime of the object

🔍 Frequent Metadata Accesses:

- Be mindful of metadata accesses along with data access for the objects

🧨 Data Delete Patterns (Early Delete Penalties)

- Ignoring delete patterns can incur early delete penalties as data tiered into lower tiers generally have a minimum retention period requirement

➡️ Ignoring Data egress charges:

- Analyzing the number of accesses alone and not amount of data read/written. (Can be 60-80% of total cost)

🗄️ Misunderstanding Data retrieval duration / charges

- Mindful of Colder-Archive tiers requiring longer retrieval times and specific charges

Microsoft

SDC 25

Key Challenges and Blind spots

Other Blind Spots

🧭 Misaligned Tiering Rules

- Having conflicting Tiering policies/rules

🧊 Ignoring Blob Size while making Tier decision

- Tiering multiple small objects to colder tiers may incur more cost in the tier change than benefit of lower storage cost

🔗 Ownership Ambiguity / Orphaned Data versions

- Missing Tiering of objects due to lack of clear ownership

🧱 Cost from the execution of the Tier change

- Ignoring the cost of executing tier change

🧬 Compliance and Audit Risks

- Be mindful of changes in location of storage / retrieval durations as they could fall under compliance and audit requirements for the user

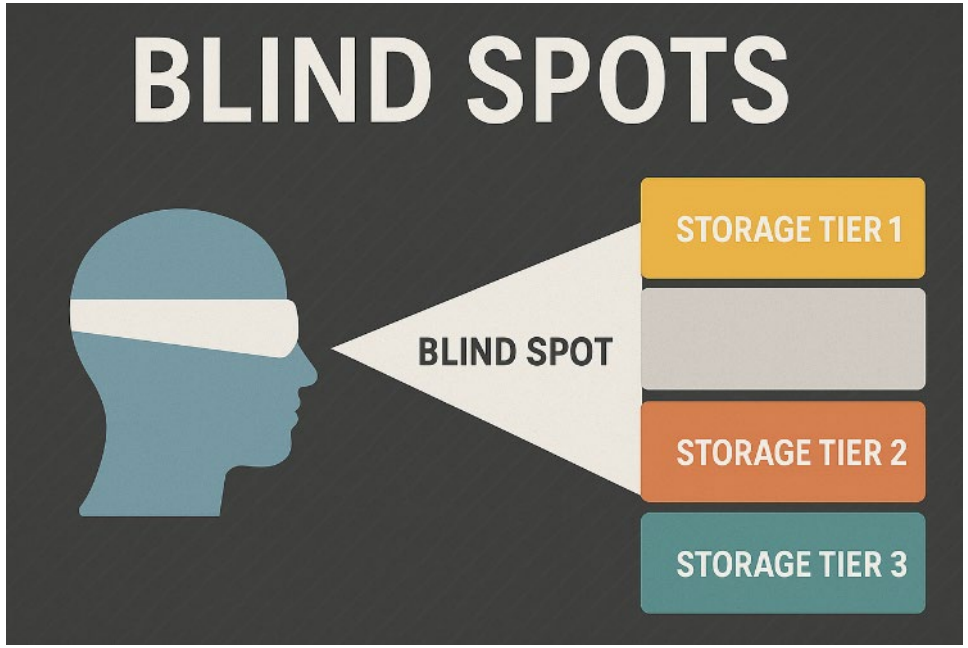

Microsoft

SDC 25

Special handling on first-time Tier Tagging

- **Policy Overlap with Retention**
  - Tiering policies may not account for retention characteristics, causing data to be moved to tiers where deletion is restricted or penalized.
  - Retention and tiering rules might conflict, especially during the first 30–90 days of tiering

- **Intelligent Tiering Override Confusion**
  - Users manually tiering data may not realize intelligent tiering still acts on those objects unless explicitly disabled, causing unexpected transitions.

- **Inaccurate High-Level Analysis**
  - Inaccurate high-level estimate of tagging since the true tagging implementation is at Blob level.
  - High Tiering execution cost in initial days since ALL data is being tiered.

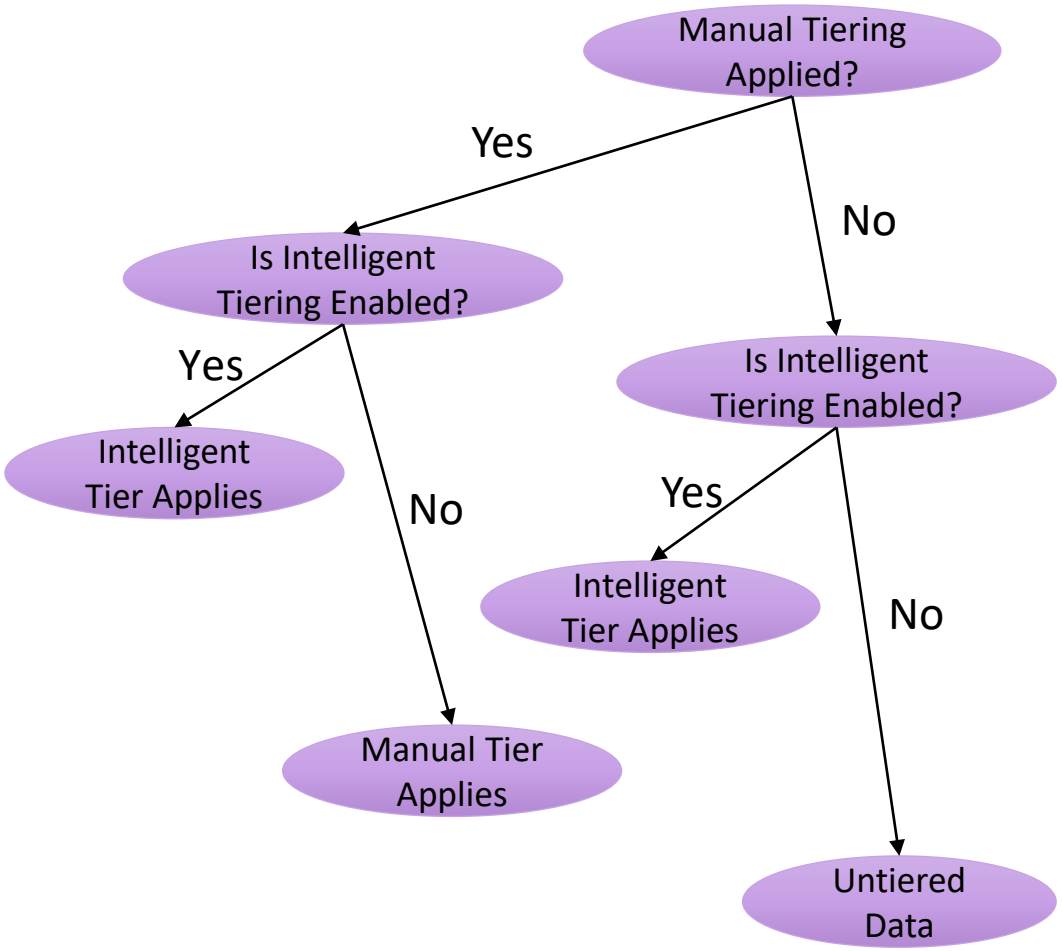

Intelligent Tiering Override Confusion

Tier Optimization for AI workloads

Importance of Tier optimization in AI Workloads

- Changing data pattern throughout the lifecycle of the Data
- Improving resource allocation & efficiency
- Extremely cost sensitive applications
- Need for long-term data management

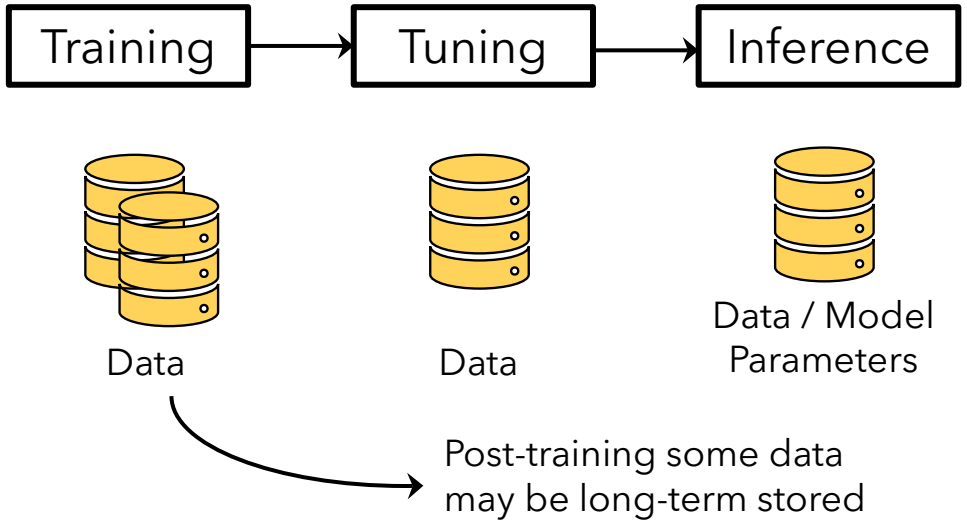

- **AI workload stages with possible Tier optimization**

	Training / Tuning	Model Inference	Long Term retained data post-training / Logs
Dataset Size	Massive	Moderate set of Model parameters	Massive
Data Usage need	Computing pipelines that are extremely sensitive to even millisecond delays	Fast access to model parameters needed along with input of new problem data being fed	Infrequently accessed until needed for another training / tuning run
Possible Tiers	Hot / Premium Tiers	Hot Tiers	Cold / Archive Tiers

Sources:
https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/scenarios/ai/infrastructure/storage
https://cloud.google.com/discover/what-is-ai-inference?hl=en
https://insideainews.com/2024/11/11/why-auto-tiering-is-essential-for-ai-solutions-optimizing-data-storage-from-training-to-long-term-archiving/

Microsoft

SDC 25

Monitoring & Visibility of Data-Usage Patterns

"You can't optimize what you don't measure."

Leverage storage analytics (AWS Storage Lens, Azure Monitor / Azure Storage Discovery, GCP Storage Insights)

- **Track metrics:** data by tier, transactions by tier, egress costs per tier, last access times, retrieval costs along with data stored cost.
- **Set alerts** for anomalies
  - Alerts for sudden spike in reads
  - Alerts to identify stale data (e.g. buckets with 80% data unused in 6+ months)
- **AI Recommendations**
  - Eg: Azure Dashboards with Copilot recommendations

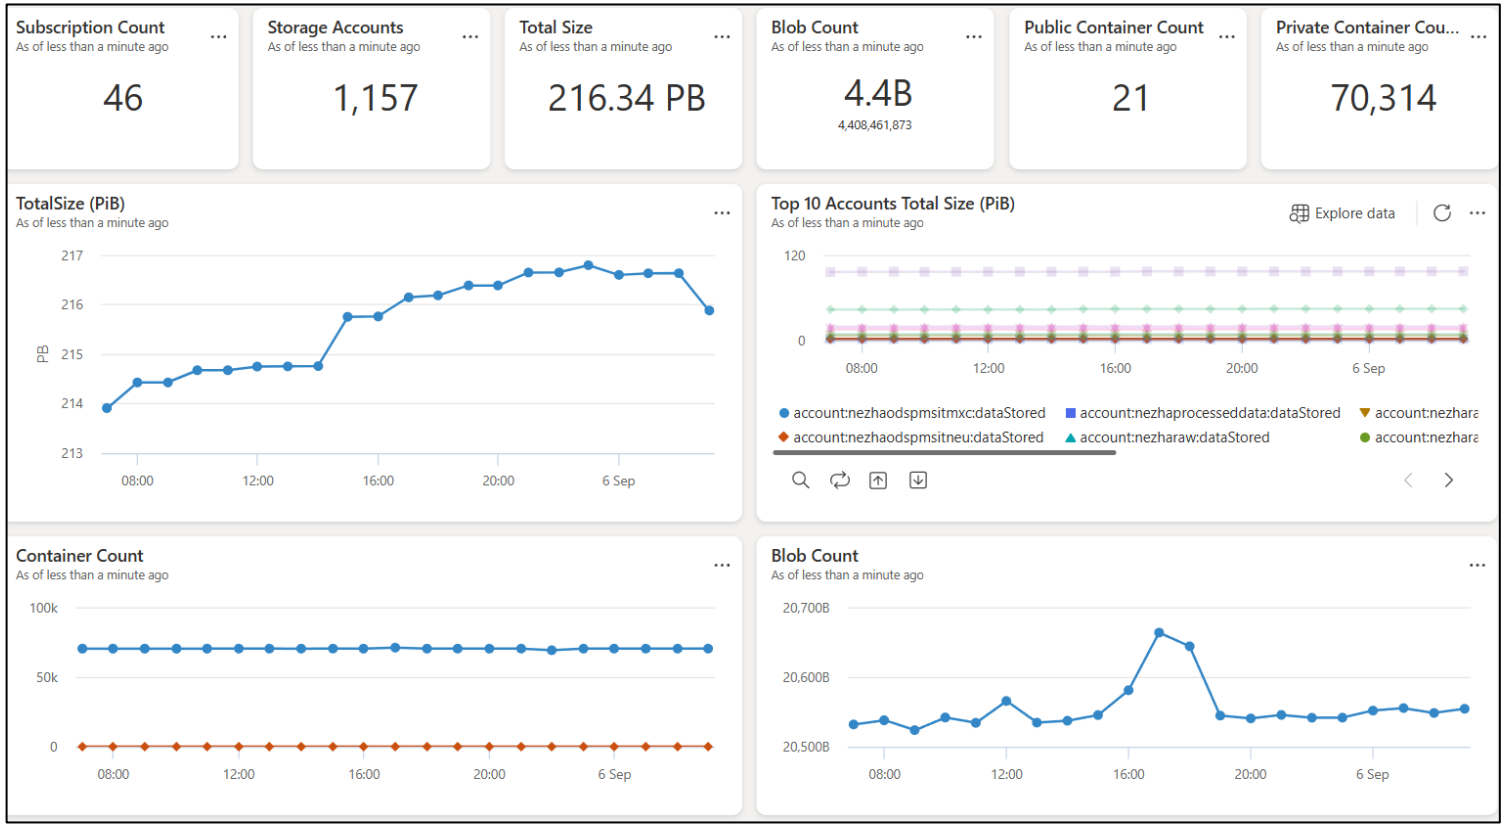

Best Practices for Cost Optimization

- Early Data Classification & Tagging
- Implement Lifecycle Actions or Intelligent Tiering
- Continuously Monitor usage and access patterns
- Hone your Retention and Deletion Practices
- Optimize Access and Cost Trade-offs
- Right-size the infra provisioning
- Apply Compression and Deduplication
- Leverage AI & Advanced Analytics
- Conduct Regular Audits

Microsoft

SDC 25

Future Tiering Innovations & Enhancements

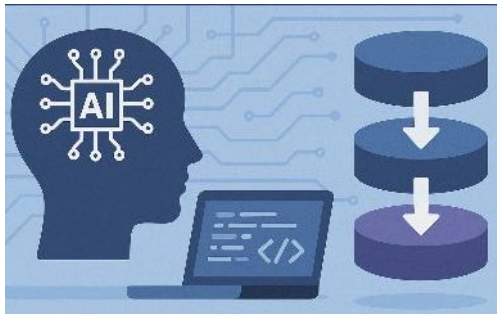

AI-First Storage Tiering

More AI-driven automation allowing systems to predict future Tiering needs

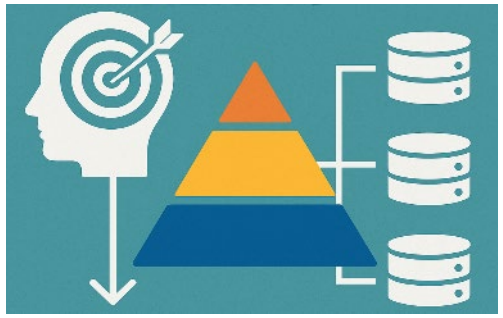

Intent-Based Tiering (V/S Policy-Based Tiering)

Minimalistic Tiering Intent needed to be setup - such as “minimize cost”, “minimize latency”

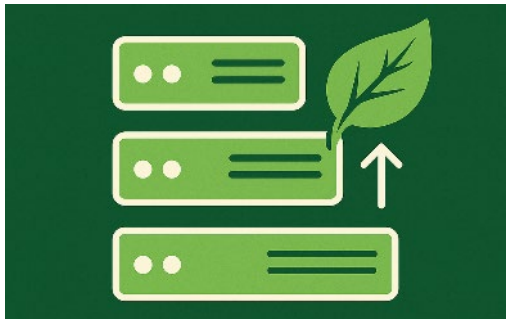

Greener Storage Tiers

Optimizing storage tiers directly to attain energy savings.

For example - Tiers based on Peak & Off-Peak time of accesses

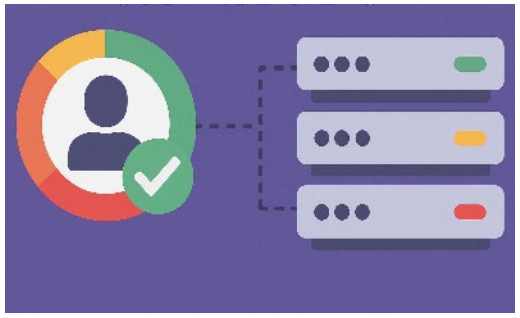

Holistic Score Based Storage Tiers

Tiers based on score assignment incorporating multiple aspects

Microsoft

SDC 25

Key Takeaways

Microsoft

SDC 25

Call To Action

Tier data and refine Tagging in a feedback loop, leveraging latest industry tools, to best optimize storage costs!

Don’t Miss out!

Modern Data Management at Exabyte Scale – With Visibility, Efficiency, and Control
Winchester
Wed Sep 17 | 1:30pm

Cloud Storage Considerations for Retrieval Augmented Generation (RAG) in AI Applications
Winchester
Wed Sep 17 | 8:30am

Microsoft

SDC 25

References

- **Azure Cost Estimation Guide:**
  - Estimate the cost of using Azure Blob Storage | Microsoft Learn
  - Estimate the cost of archiving data (Azure Blob Storage) - Azure Storage | Microsoft Learn
- **Azure Blob Lifecycle Management**
  - Azure Blob Storage lifecycle management overview - Azure Blob Storage | Microsoft Learn
- **Azure Storage Actions**
  - About Azure Storage Actions - Azure Storage Actions | Microsoft Learn
- **Azure Storage Discovery**
  - Azure Storage Discovery Overview | Microsoft Learn
- **Cost Optimization Success Stories:**
  - Unlock seamless data management with Azure Storage Actions—now generally available | Microsoft Azure Blog

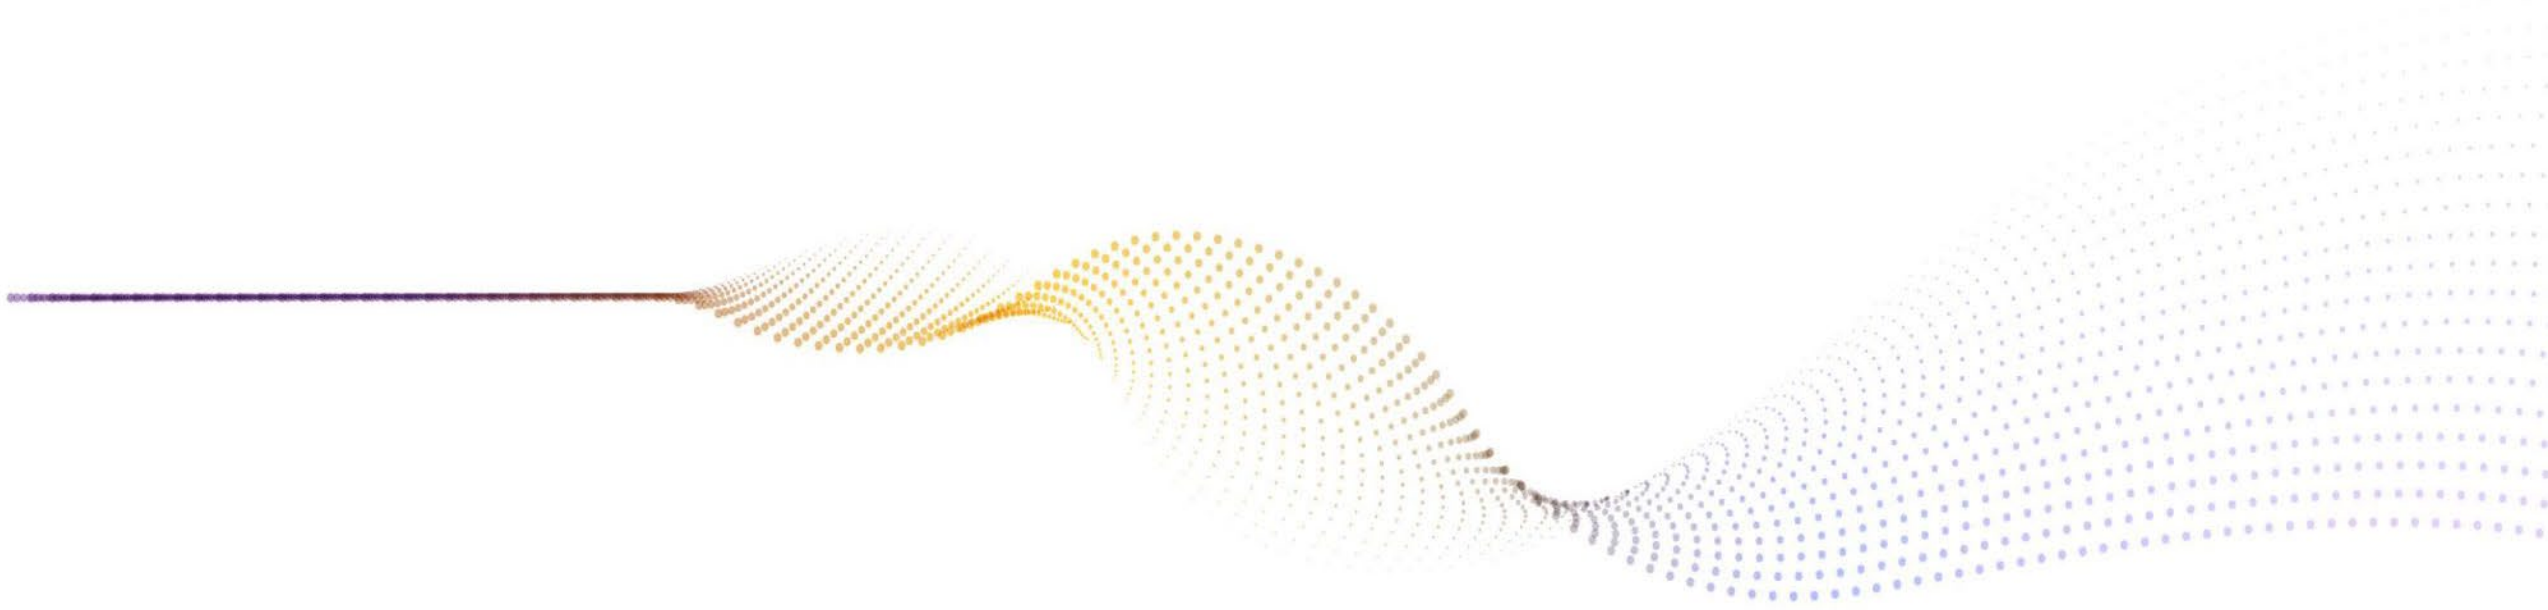

Thank you for attending!

Please remember to rate this session. You get access the presentations at
http://sniadeveloper.org/conference

Microsoft

SDC 25

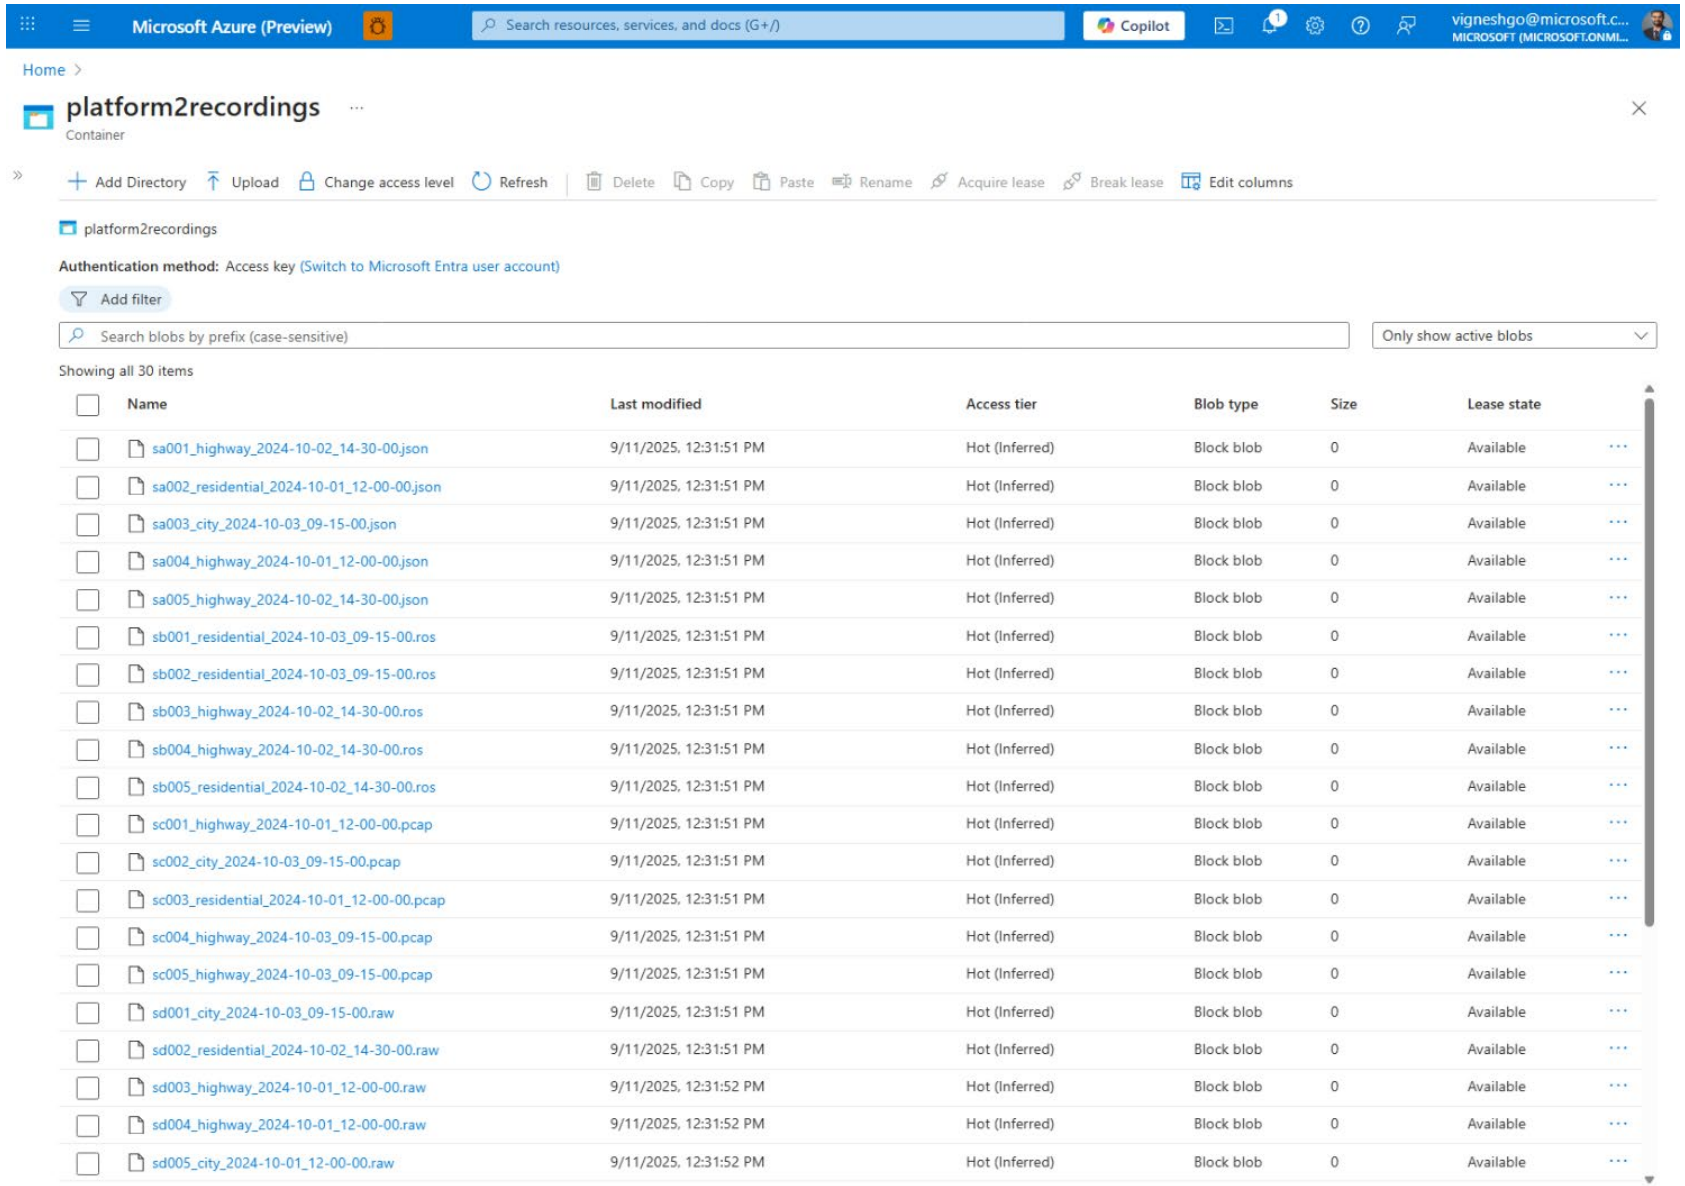

Microsoft

SDC 25

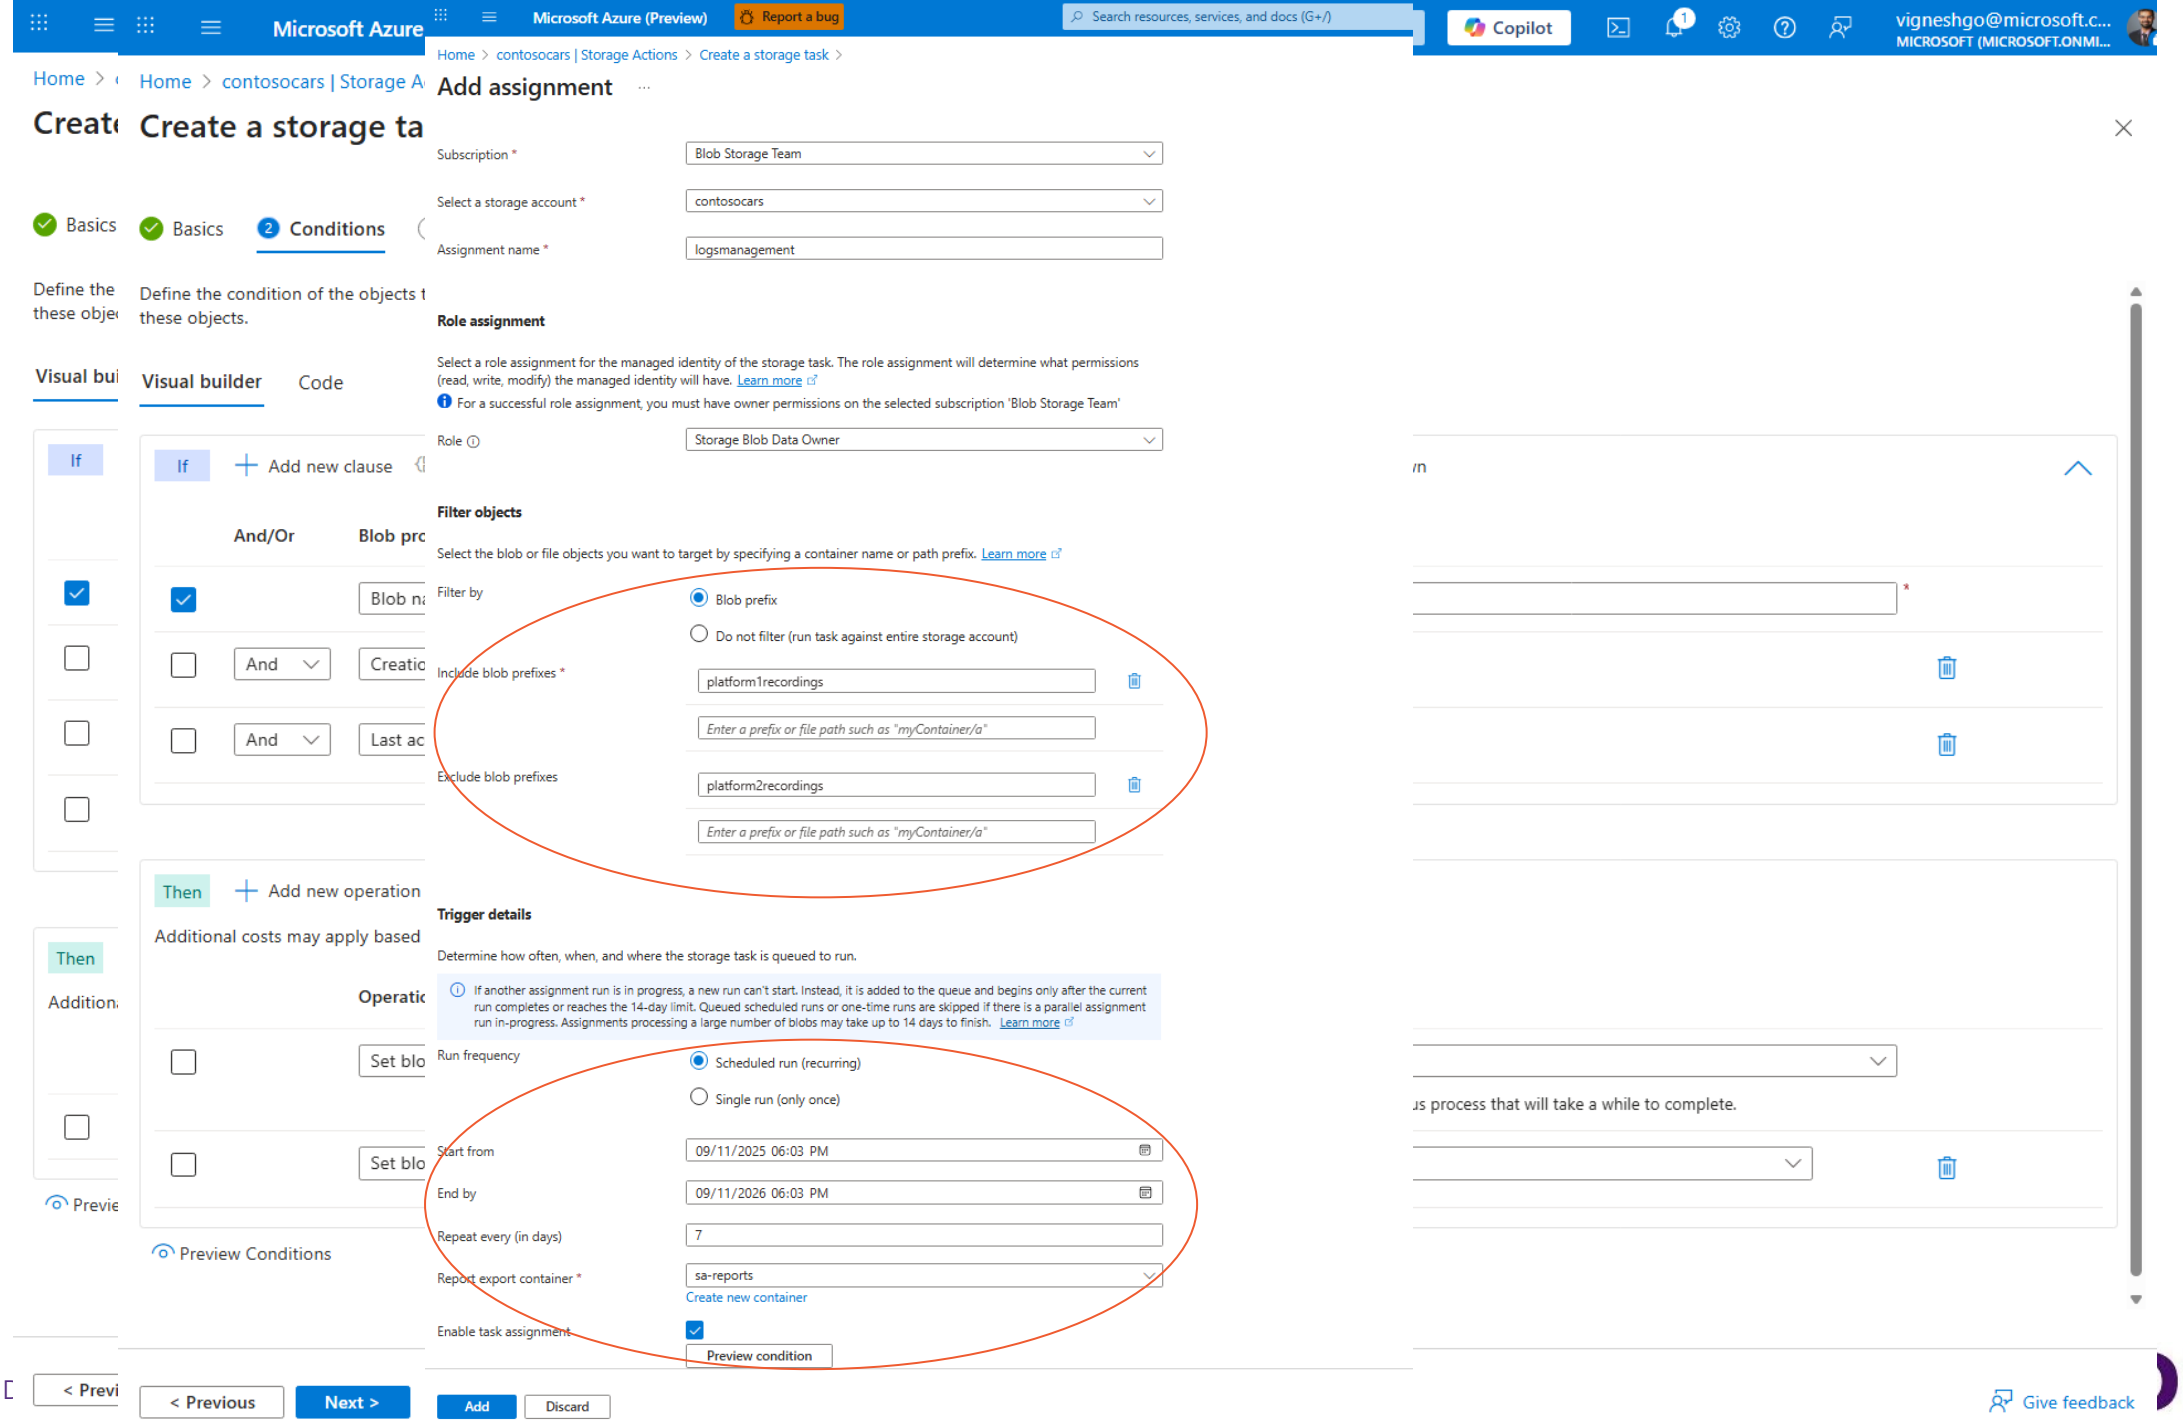
Microsoft Azure (Preview)
Report a bug
Search resources, services, and docs (G+/)
Copilot
vigneshgo@microsoft.c...
MICROSOFT (MICROSOFT.ONMI...
Home > contosocars | Storage Actions > Create a storage task >
Add assignment
Subscription *
Blob Storage Team
Select a storage account *
contosocars
Assignment name *
logsmanagement
Role assignment
Select a role assignment for the managed identity of the storage task. The role assignment will determine what permissions (read, write, modify) the managed identity will have. Learn more
For a successful role assignment, you must have owner permissions on the selected subscription 'Blob Storage Team'
Role
Storage Blob Data Owner
Filter objects
Select the blob or file objects you want to target by specifying a container name or path prefix. Learn more
Filter by
Blob prefix
Do not filter (run task against entire storage account)
Include blob prefixes *
platform1recordings
Enter a prefix or file path such as "myContainer/a"
Exclude blob prefixes
platform2recordings
Enter a prefix or file path such as "myContainer/a"
Trigger details
Determine how often, when, and where the storage task is queued to run.
If another assignment run is in progress, a new run can't start. Instead, it is added to the queue and begins only after the current run completes or reaches the 14-day limit. Queued scheduled runs or one-time runs are skipped if there is a parallel assignment run in-progress. Assignments processing a large number of blobs may take up to 14 days to finish. Learn more
Run frequency
Scheduled run (recurring)
Single run (only once)
Start from
09/11/2025 06:03 PM
End by
09/11/2026 06:03 PM
Repeat every (in days)
7
Report export container *
sa-reports
Create new container
Enable task assignment
Preview condition
Add
Discard
Home > contosocars | Storage A
Create a storage ta
Basics
Basics
2 Conditions
Define the condition of the objects t
these objects.
Visual builder
Code
If
Add new clause
And/Or
And
And
Then
Add new operation
Additional costs may apply based
Preview Conditions
< Previous
Next >
Give feedback

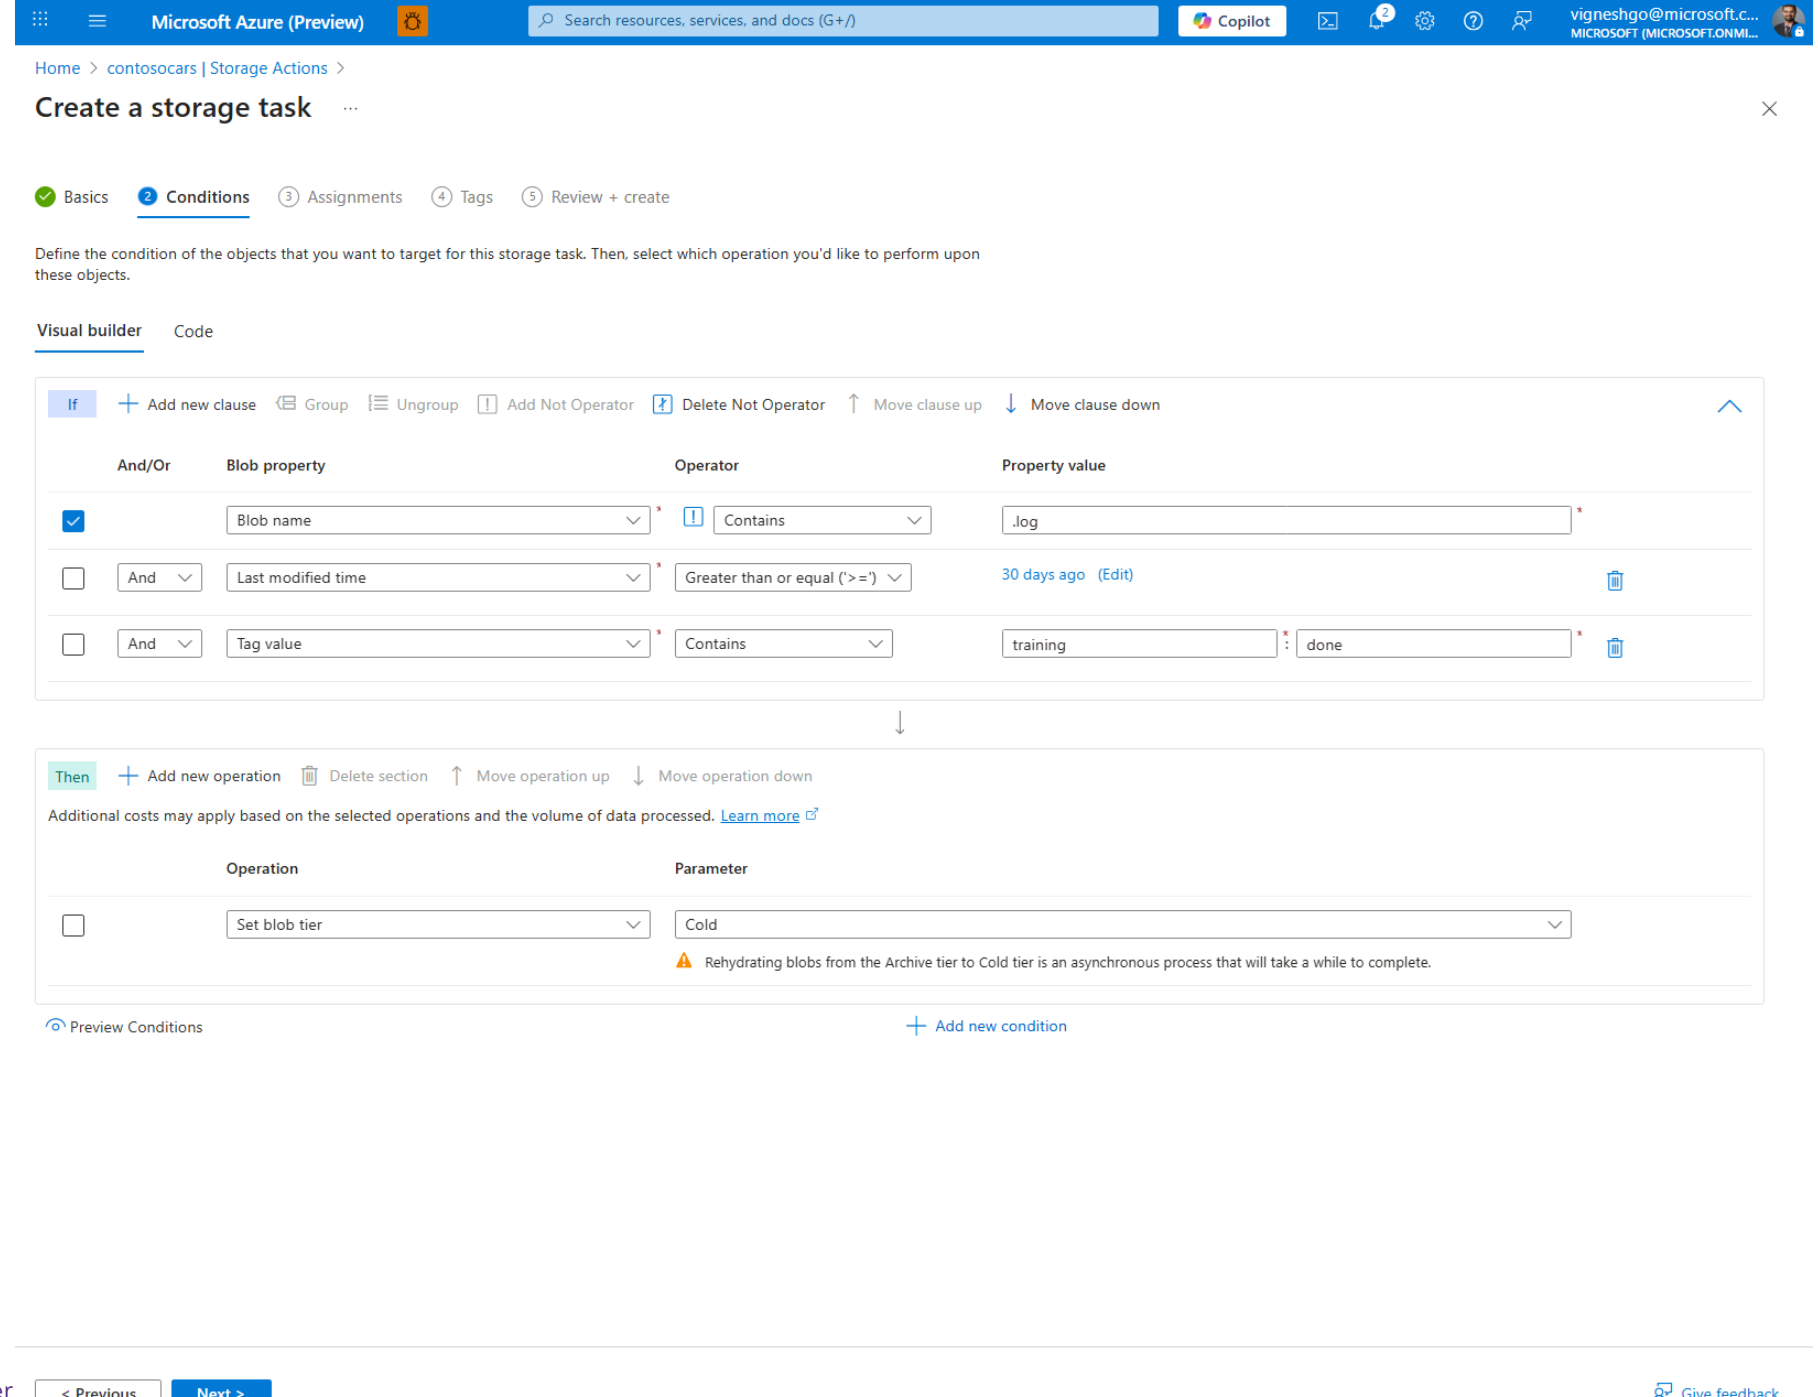
Microsoft Azure (Preview)
Search resources, services, and docs (G+/)
Copilot
vigneshgo@microsoft.c...
MICROSOFT (MICROSOFT.ONMI...
Home > contosocars | Storage Actions >
Create a storage task
Basics
2 Conditions
3 Assignments
4 Tags
5 Review + create
Define the condition of the objects that you want to target for this storage task. Then, select which operation you'd like to perform upon these objects.
Visual builder
Code
If
Add new clause
Group
Ungroup
Add Not Operator
Delete Not Operator
Move clause up
Move clause down
And/Or
Blob property
Operator
Property value
Blob name
Contains
.log
And
Last modified time
Greater than or equal ('>=')
30 days ago (Edit)
And
Tag value
Contains
training
done
Then
Add new operation
Delete section
Move operation up
Move operation down
Additional costs may apply based on the selected operations and the volume of data processed. Learn more
Operation
Parameter
Set blob tier
Cold
Rehydrating blobs from the Archive tier to Cold tier is an asynchronous process that will take a while to complete.
Preview Conditions
Add new condition
< Previous
Next >
Give feedback