
Unlocking Cost Efficiency: Reserved Instances and Savings Plans for Higher Education

Gotthard Szabo
Adam Pritchard
August 1, 2024

Housekeeping



If you have any technical issues/questions, please message Claire directly in the chat.



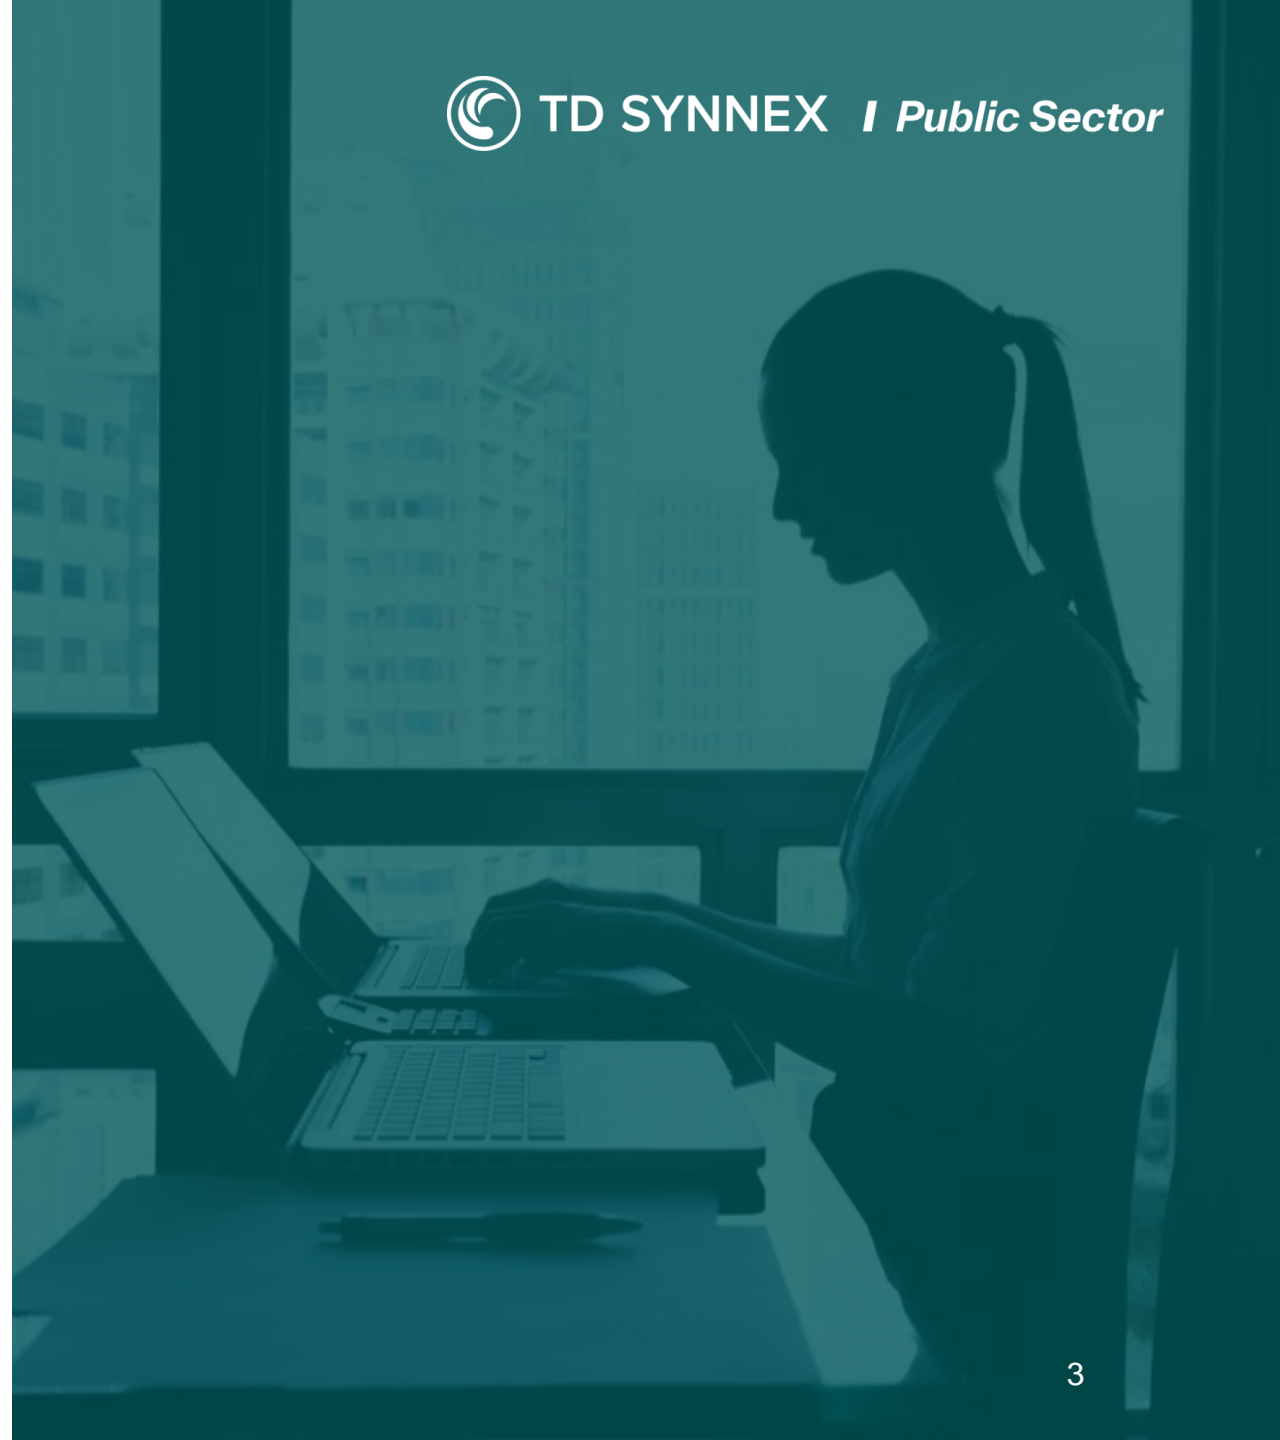
All lines are MUTED for the duration of the webinar, however, if you have questions throughout the presentation, feel free to drop them in the Q&A or chat field.



This session is being recorded. The recording of the webinar will be made available 1 week after the live session.

Agenda

- Contract Overview
- Introduction to FinOps
- Reserved Instances
- Savings Plans
- Q&A



Contract Details

Contract Name:	• Internet2
Scope:	• AWS & Professional Services
Customer Category:	• Higher Education
Period of Performance:	• 9/30/2014 - 5/12/2025 **Auto renews each year consecutively following 2025**
Authorized Dealer Permitted:	• Yes

- Contract Standard Discount - 5%
- NIH STRIDES - 12% discount, Free AWS Enterprise Support
- AWS Resold Training – 5%
- Data Egress Fee waiver
- Immersion Day Training at No cost
- Well Architected Framework Reviews at No cost
- FinOps services at No cost
- Access to Professional Services and Managed Services
- **Optional** AWS Partner Led Business Support
- FERPA compliance
- DLT Custom Portal for I2 Members

Today's Subject Matter Expert

- Gotthard Szabo
 - Sr. Director, Cloud Operations



Introduction

- Commitment based discounts should be only a part of a mature cost optimization strategy
- There are no restrictions for TD SYNEX customers to implement their own cost optimization strategy, including purchasing SPs and RIs
- TD SYNEX can help you with all six pillars of AWS Well-Architected Framework
- TD SYNEX can implement and manage commitment-based cost optimization for your accounts
- Risk adjusted returns for self-service commitment-based cost optimization can be negative or low
- Risk adjusted returns for TD SYNEX managed commitment-based cost optimization can be high
- TD SYNEX Ranked #71 on the 2024 FORTUNE 500
- One of the largest cloud services distributors worldwide
- We have the experience and resources

AWS Well-Architected Framework

- The framework is based on six pillars.
 1. Operational Excellence - running and monitoring systems
 2. Security - protecting information and systems
 3. Reliability - performing intended functions, recover quickly from failure
 4. Performance Efficiency - structured and streamlined allocation of IT and computing resources
 - 5. Cost Optimization - avoiding unnecessary costs**
 6. Sustainability - minimizing environmental impacts of running cloud workloads

The focus of this talk is cost optimization.

AWS WAF, Cost Optimization

There are five focus areas for cost optimization in the cloud.

1. Practice Cloud Financial Management
2. Expenditure and usage awareness
3. Cost-effective resources
4. Manage demand and supplying resources
5. Optimize over time

AWS WAF, CO, Cloud Financial Management

- Establish ownership of cost optimization
- Establish a partnership between finance and technology
- Establish cloud budgets and forecasts
- Implement cost awareness in your organizational processes
- Report and notify on cost optimization
- Monitor cost proactively
- Keep up-to-date with new service releases
- Create a cost-aware culture
- Quantify business value from cost optimization

AWS WAF, CO, Expenditure and usage awareness

- Develop policies based on your organization requirements
- Implement goals and targets
- Implement an account structure
- Implement groups and roles
- Implement cost controls
- Track project lifecycle
- Identify cost attribution categories
- Establish organization metrics
- Configure billing and cost management tools
- Track resources over their lifetime
- Implement a decommissioning process
- Decommission resources
- Decommission resources automatically
- Enforce data retention policies

AWS WAF, CO, Cost Effective Resources

- Identify organization requirements for cost
- Analyze all components of the workload
- Perform a thorough analysis of each component
- Select software with cost-effective licensing
- Select components of this workload to optimize cost in line with organization priorities
- Perform cost analysis for different usage over time

- Select the correct resource type, size, and number
- Perform cost modeling
- Select resource type, size, and number based on data
- Select resource type, size, and number automatically based on metrics
 - AWS Auto Scaling
 - Rightsizing with AWS Compute Optimizer and Memory Utilization Enabled
 - Amazon S3 Storage Lens
 - Amazon S3 Intelligent Tiering and Amazon EFS Infrequent Access

AWS WAF, CO, Cost Effective Resources

Only NOW we come to selecting the best pricing model! Lot to be done till this point!

- On-Demand Instances
- Spot Instances
- **Commitment discounts - Savings Plans**
- **Commitment discounts - Reserved Instances**

- Select third-party agreements with cost-efficient terms
- Implement pricing models for all components of this workload
- Perform pricing model analysis at the management account level

AWS WAF, CO, Cost Effective Resources

- Choose Regions based on cost
 - Perform data transfer modeling
 - Architect based on data transfer needs
 - Select components to optimize data transfer cost
 - Implement services to reduce data transfer costs
-
- AWS DEFW program – data egress fee waiver
 - TDS CloudFront PPA – up to 80% discount

AWS WAF, Cost Optimization

The remaining two pillars:

Manage demand and supply resources

- Perform an analysis on the workload demand
- Implement a buffer or throttle to manage demand
- Supply resources dynamically

Optimize over time

- Develop a workload review process
- Review and analyze this workload regularly

Commitment Discounts - Savings Plans

Savings Plans can reduce your bill by up to 72% compared to On-Demand prices

One- or three-year hourly spend commitment.

AWS offers three types of Savings Plans:

1. Compute Savings Plans
2. EC2 Instance Savings Plans
3. Amazon SageMaker Savings Plans

Commitment Discounts - Savings Plans

3 different payment options

1. No Upfront option does not require any upfront payment, and your commitment will be charged purely on a monthly basis.

2. Partial Upfront option offers lower prices on Savings Plans. You are charged at least half of your commitment upfront, and the remaining will be charged on a monthly basis.

3. All Upfront option, you will receive the lowest prices, and your entire commitment will be charged in one payment.

Commitment Discounts - Compute Savings Plans

- Provide the most flexibility
- Can reduce your costs by up to 66%
- One- or three-year terms
- Commitment to a consistent amount of usage (\$/hour)
- Plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, region, OS or tenancy
- Plans also apply to Fargate and Lambda usage

Commitment Discounts – EC2 Savings Plans

- Provide the lowest prices
- Savings up to 72%
- One- or three-year terms
- Commitment to a consistent amount of usage (\$/hour) of individual instance families in a region

Commitment Discounts – Amazon SageMaker SPs

- Can reduce your costs by up to 64%
- One- or three-year terms
- Commitment to a consistent amount of usage (\$/hour)
- Flexible pricing model for Amazon SageMaker
- Plans automatically apply to eligible SageMaker services regardless of instance family, size, or region

Commitment Discounts - Savings Plans

- Savings Plans can be purchased in any account within an AWS Organization/Consolidated Billing family.
- Savings Plans benefits can be restricted to the account that purchased the SPs.
- SPs benefits can be shared with the entire consolidated billing family of accounts.
- SPs benefits cannot be shared amongst a subset of billing family accounts.
- Consolidated Billing Family, sharing disabled - Savings Plans are applied only to the owner account's usage.
- Consolidated Billing Family, sharing enabled - Savings Plans are applied first to the owner account's usage, and then to other accounts' usage.

Commitment Discounts - Savings Plans

- Savings Plans apply automatically to eligible AWS usage.
- TD SYNEX doesn't override AWS application logic.
- Amazon EC2 Reserved Instances are applied before Savings Plans.)
- EC2 Savings Plans are applied before Compute Savings Plans.
- Savings Plans are applied to your highest savings percentage first.
- For multiple equal savings percentages, Savings Plans are applied to the first usage with the lowest Savings Plans rate.
- Savings Plans continue to apply until there are no more remaining usages, or your commitment is exhausted.

Commitment discounts - Savings Plans Risks

Not returnable unless:

- Has an hourly commitment of \$100 or less
- Purchased in the past 7 days
- Purchased in the same calendar month as when the request is made

Commitment discounts - Savings Plans Risks

Overcommitment – you pay your hourly commitment for one or three years whether you use services or not.

Undercommitment – loss of savings opportunity.



Commitment Discounts - Reserved Instances

- Reserved Instances (RIs) are not physical instances, but rather a billing discount.
- Reserved Instances apply automatically to eligible AWS usage.
- With RIs, you make a commitment to use a number of various EC2 instances during the reservation period.
- AWS offers reservations for the following services, each with their own restrictions.

AWS EC2

AWS RDS

AWS OpenSearch Service

Amazon Redshift

Amazon ElastiCache

Amazon DynamoDB


Commitment Discounts - EC2 RI Instances

- Instance attributes

A Reserved Instance has four instance attributes that determine its price:

1. **Instance type:** For example, m4.large. This is composed of the instance family (for example, m4) and the instance size (for example, large).
2. **Region:** The Region in which the Reserved Instance is purchased.
3. **Tenancy:** Whether your instance runs on shared (default) or single-tenant (dedicated) hardware.
4. **Platform:** The operating system; for example, Windows or Linux/Unix.

Commitment Discounts - EC2 Instance types



Commitment Discounts - EC2 RI Instance attributes

Term commitment

You can purchase a Reserved Instance for a one-year or three-year commitment, with the three-year commitment offering a bigger discount.

- **One-year:** A year is defined as 31536000 seconds (365 days).
- **Three-year:** Three years is defined as 94608000 seconds (1095 days).

Risk – your usage pattern might change significantly during the term!

Commitment Discounts - EC2 RI Instances

RI Payment options

- **All Upfront:** Full payment is made at the start of the term, with no other costs or additional hourly charges incurred for the remainder of the term, regardless of hours used.
- **Partial Upfront:** A portion of the cost (usu. 50% or more) must be paid upfront and the remaining hours in the term are billed at a discounted hourly rate, regardless of whether the Reserved Instance is being used.
- **No Upfront:** You are billed a discounted hourly rate for every hour within the term, regardless of whether the Reserved Instance is being used. No upfront payment is required.

Regional and Zonal Reserved Instances (scope)

Regional: When you purchase a Reserved Instance for a Region, it's referred to as a regional Reserved Instance.

Zonal: When you purchase a Reserved Instance for a specific Availability Zone, it's referred to as a zonal Reserved Instance.

Regional and zonal Reserved Instances (scope)

	Regional Reserved Instances	Zonal Reserved Instances
Ability to reserve capacity	A regional Reserved Instance does <i>not</i> reserve capacity.	A zonal Reserved Instance reserves capacity in the specified Availability Zone.
Availability Zone flexibility	The Reserved Instance discount applies to instance usage in any Availability Zone in the specified Region.	No Availability Zone flexibility—the Reserved Instance discount applies to instance usage in the specified Availability Zone only.
Instance size flexibility	The Reserved Instance discount applies to instance usage within the instance family, regardless of size. Only supported on Amazon Linux/Unix Reserved Instances with default tenancy.	No instance size flexibility—the Reserved Instance discount applies to instance usage for the specified instance type and size only.
Queuing a purchase	You can queue purchases for regional Reserved Instances.	You can't queue purchases for zonal Reserved Instances.

Commitment Discounts - EC2 RI Instances

- Offering classes – Standard and Convertible
- If your computing needs change, you might be able to modify or exchange your Reserved Instances, depending on the offering class.

Standard and Convertible EC2 RI Differences

	Standard Reserved Instance	Convertible Reserved Instance
Modify Reserved Instances	Some attributes can be modified.	Some attributes can be modified.
Exchange Reserved Instances	Can't be exchanged.	Can be exchanged during the term for another Convertible Reserved Instance with new attributes, including instance family, instance type, platform, scope, or tenancy.
Sell in the Reserved Instance Marketplace	Can be sold in the Reserved Instance Marketplace.	Can't be sold in the Reserved Instance Marketplace.
Buy in the Reserved Instance Marketplace	Can be bought in the Reserved Instance Marketplace.	Can't be bought in the Reserved Instance Marketplace.

Modify EC2 Reserved Instances

Attributes	Subset	Merge	Split	Frequency
You can modify attributes such as the Availability Zone, instance size (within the same instance family and generation), and scope of your Reserved Instance.	You can modify all or a subset of your Reserved Instances.	You can merge two or more Reserved Instances into a single Reserved Instance.	You can split your Reserved Instances.	You can modify your reservations as frequently as you like.

How Reserved Instances are applied

- Reserved Instances are not physical instances, but rather a billing discount that is applied to the running On-Demand Instances in your account.
- Reserved Instances apply automatically to eligible AWS usage.
- Amazon EC2 Reserved Instances are applied before Savings Plans.
- The offering class (Standard or Convertible) of the Reserved Instance does not affect how the billing discount is applied.

RI marketplace restrictions and limitations

- Only Amazon EC2 Standard regional and zonal Reserved Instances can be sold in the Reserved Instance Marketplace.
- Amazon EC2 Convertible Reserved Instances cannot be sold in the Reserved Instance Marketplace.
- Reserved Instances for other AWS services, such as Amazon RDS and Amazon ElastiCache, cannot be sold in the Reserved Instance Marketplace.
- There must be at least one month remaining in the term of the Standard Reserved Instance.
- You cannot sell discounted SRIs on the Marketplace.
- You can sell No Upfront, Partial Upfront, or All Upfront Reserved Instances in the Reserved Instance Marketplace as long as they have been active in your account for at least 30 days.
- If there is an upfront payment on a Reserved Instance, it can only be sold after AWS has received the upfront payment.
- AWS charges a service fee of 12 percent of the seller price of each Standard Reserved Instance you sell in the Reserved Instance Marketplace.

RI sharing

- Amazon EC2 Reserved Instance benefits are applied before Savings Plans.
- RI benefits can be restricted to the account that purchased the RIs.
- RI benefits can be shared with the entire consolidated billing family of accounts.
- Consolidated Billing Family of an Organization
 - Sharing disabled – RIs are applied only to the owner account's usage.
 - Sharing enabled – RIs are applied first to the owner account's usage, and then to other accounts' usage. Both owner and receiving accounts must have sharing enabled set.

RI Usage billing

- Reserved Instances are billed for every clock-hour during the term that you select, regardless of whether an instance is running.
- A Reserved Instance billing benefit can apply to a maximum of 3600 seconds (one hour) of instance usage per clock-hour.

	1:00	1:15	1:30	1:45
Instance 1				
Instance 2				
Instance 3				
Instance 4				

	1:00	1:15	1:30	1:45
Instance 1				
Instance 2				
Instance 3				
Instance 4				

Uses Reserved Instance Rate for first 3600 seconds of use
Uses On-Demand Rate

Liability, Investments – RIs and Savings Plans

- Upfront purchases are expensive, can cause cash-flow issues
- Partial-upfront and no-upfront purchases also have significant liabilities associated with them
- Even a mid-size, \$100/hr, commitment equates to
 - \$876,000 liability for one year, \$2,628,000 for three years
- TD SYNEX manages many thousands of dollars in hourly commitments, tens of millions of dollars in liabilities

Convertible RIs

- Commitment - the post-discount dollar amount of discount instruments held that offset a certain amount of usage. For instance, \$1 per hour of commitment that's getting a 50% discount on a given instance will offset \$2 per hour of usage.
- Instances will have expiration dates in addition to other attributes (e.g. term, prepayment preference). Commitment is paid to AWS regardless of whether or not it's offsetting actual usage – all the way until the commitment expires.
- Total commitment amount cannot change, but we can reduce per-hour commitments by merging all or part of commitments with commitments having later expiration dates.

Convertible RIs

- If you merge two or more Convertible Reserved Instances, the term of the new Convertible Reserved Instance must be the same as the old Convertible Reserved Instances, or the highest of the Convertible Reserved Instances.
- The expiration date for the new Convertible Reserved Instance is the expiration date that's furthest in the future.

Reserved Instance ID	Term	Expiration date
aaaa1111	1-year	2024-12-31
bbbb2222	1-year	2024-07-31
cccc3333	3-year	2025-06-30
dddd4444	3-year	2026-12-31

RI and SP Coverage Goal



RI and SP Risk Scenarios

Are you ready to handle risky situations?

- Purchase was made for the wrong class
- EC2 class requirements change
- Architectural changes, transition to serverless computing
- Workload moves to another hyperscaler
- Purchase was made for the wrong region
- Customer moves their workload from one region to another region
- Regional workload doubles
- Regional workload drops by 50%
- Customer turns off their workloads Saturday and Sunday
- Customer turns off their workloads between 6 PM and 6 AM
- Grant ends in less than a year

Working with TD SYNEX

- Customer purchases Savings Plans and Reserved instances
- Customer manages them, and deals with AWS in case of problems

For the customer:

- Risk level – high
- Liability – high
- Investment requirement (upfront) – high
- Time to manage – high
- Staffing - expensive
- Benefits – can vary, function of skills, and effort invested
- Risk adjusted return – negative to medium

Working with TD SYNEX

- TD SYNEX purchases Savings Plans and Reserved instances on behalf of customer
- TD SYNEX manages everything, and deals with AWS in case of problems
- TD SYNEX shares a portion of benefits with customer (TD SYNEX keeps cost of risk)

For the customer:

- Risk level – zero
- Liability – zero
- Investment requirement (upfront) – zero
- Time to manage – minimal
- Staffing - minimal
- Benefits – similar to, or higher than self-managed
- Risk adjusted return - high

Working with TD SYNEX

Benefits can be shared as:

- Fixed, or flexible overall discount rate increase for all utilizations
- Fixed, or flexible discount rate increase for arbitrageable utilization
- Monthly fixed credit posted to invoice
- Monthly varying credit, percentage of savings after cost, posted to invoice

Requires signing basic contract

Conclusion

- Commitment based discounts should be only a part of a mature cost optimization strategy
- There are no restrictions for TD SYNEX customers to implement their own cost optimization strategy, including purchasing SPs and RIs
- TD SYNEX can help you with all six pillars of AWS Well-Architected Framework
- TD SYNEX can implement and manage commitment-based cost optimization for your accounts
- Risk adjusted returns for self-service commitment-based cost optimization can be negative or low
- Risk adjusted returns for TD SYNEX managed commitment-based cost optimization can be high
- TD SYNEX Ranked #71 on the 2024 FORTUNE 500
- One of the largest cloud services distributors worldwide
- We have the experience and resources



-
- **Contact Us: Internet2NetPlus@dlt.com**
 - **Please stay with us at the end of the webinar for a quick survey that will only take about one minute**

THANK YOU!