Understanding Cloud Disaster Recovery Costs

Natalie Serebryakova
Staff Cloud Engineer

In-n-out.cloud

About me

- My name is Natalie
- I'm a Staff Cloud Engineer
- Worked at Symantec, vevo, reddit, HashiCorp

- Consulting https://www.in-n-out.cloud 
- AWS Community Builder
- LinkedIn: https://www.linkedin.com/in/login-id/

Agenda

- What is Disaster Recovery
- What is Cost Estimation and Optimization
- Business Continuity Planning
- Estimating DR Costs for Cloud resources
- Example of Cost Optimization for Services hosted on AWS Cloud (EKS, RDS)

* examples and financial numbers used in this presentation are not the real data

What is a Disaster?

Natural Disaster	Technical Failures	Human Action
Earthquake	Power failure	Unintentional miscofiguraiton
Floods	Network connectivity	Unauthorized access or modifications

High Availability

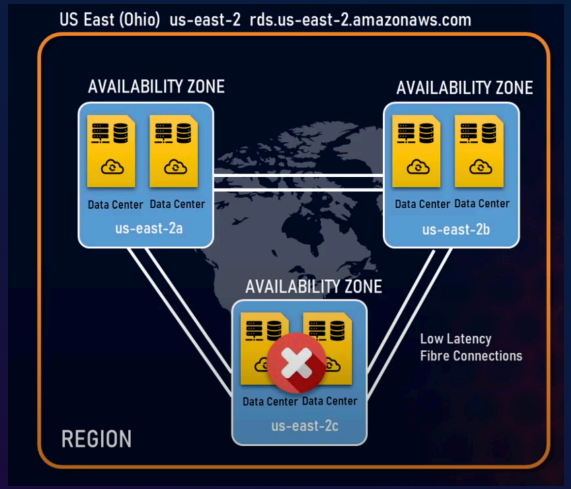

High Availability

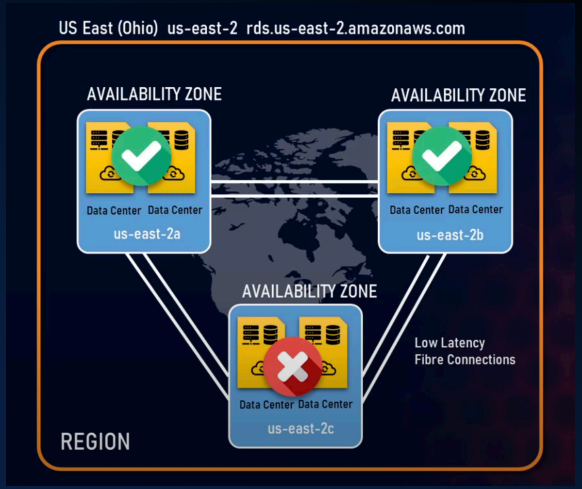

Disaster Recovery

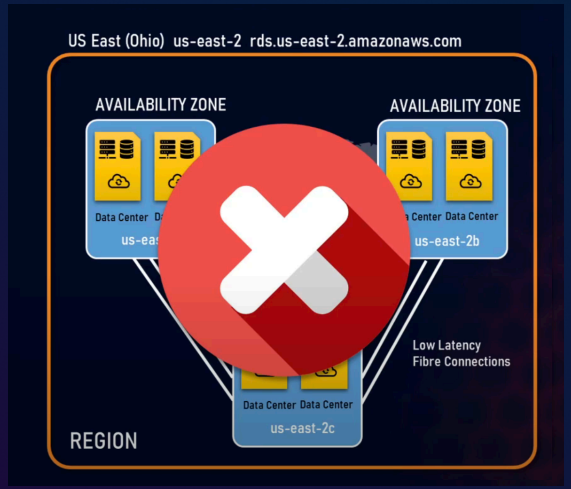

Disaster Recovery

High Availability is not a Disaster Recovery

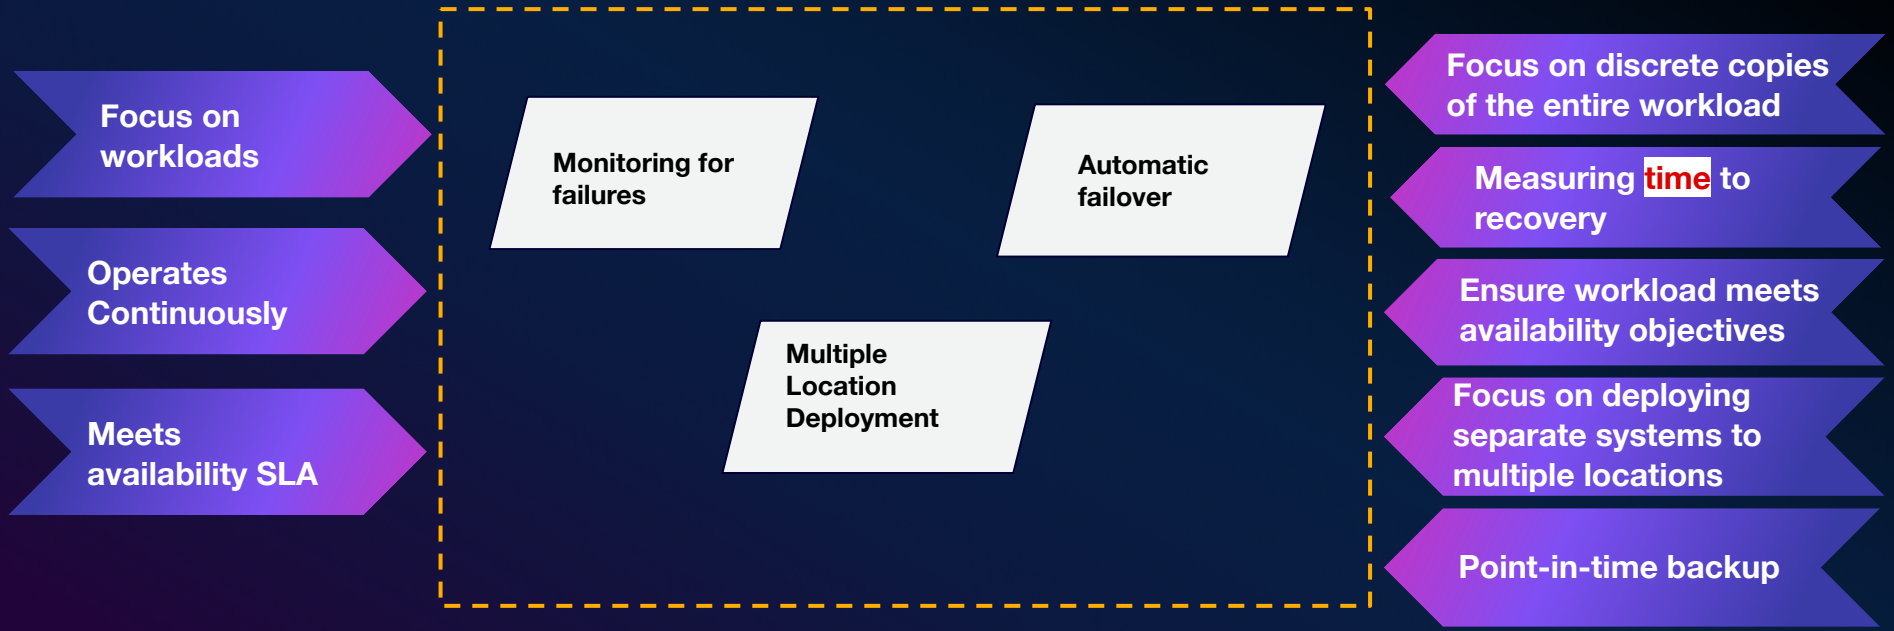

Resiliency: workload ability to recover quickly

DR and Availability

DR measures objectives for one-time events

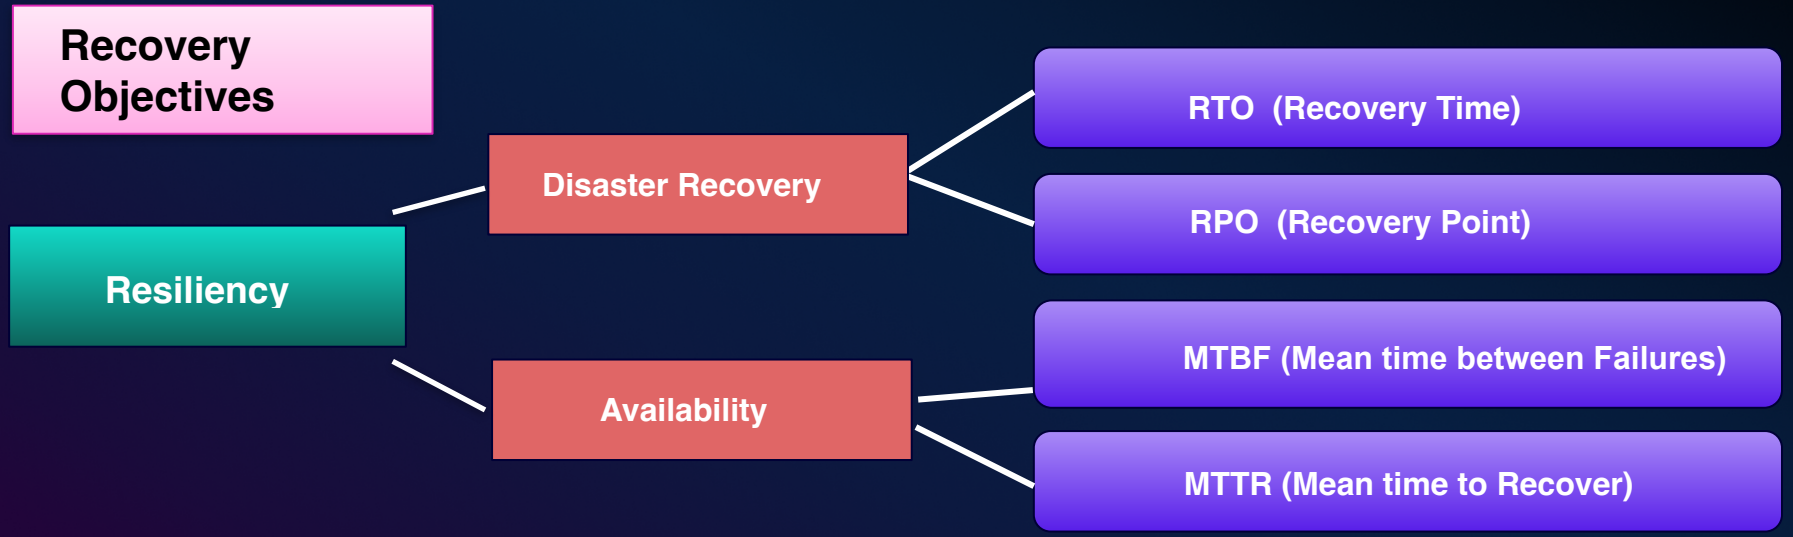

RPO and RTO

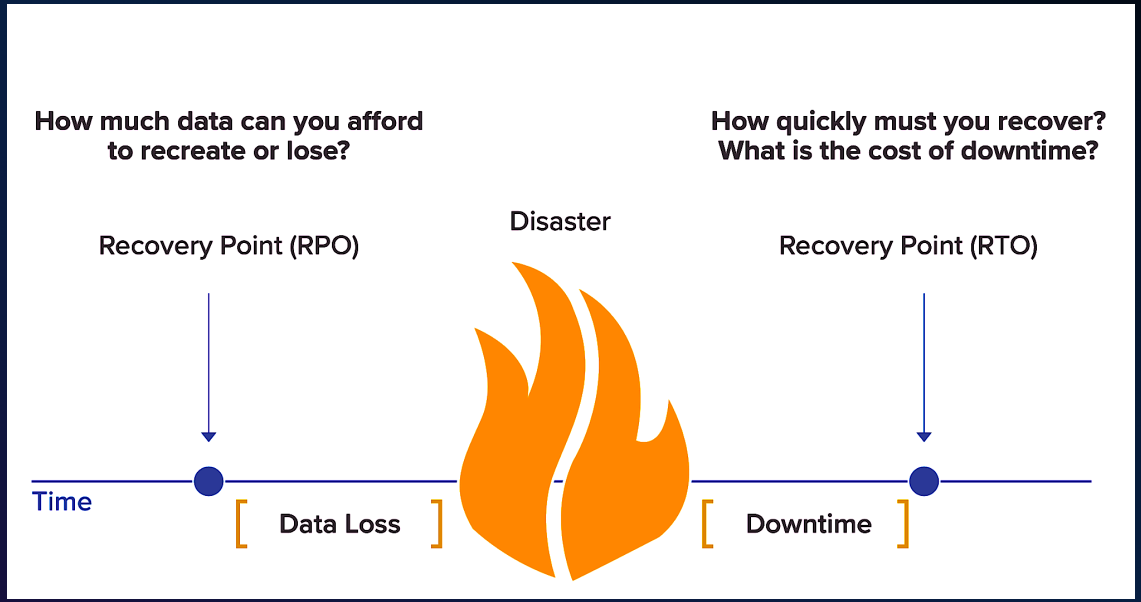
How much data can you afford to recreate or lose?
How quickly must you recover? What is the cost of downtime?
Disaster
Recovery Point (RPO)
Recovery Point (RTO)
Time
Data Loss
Downtime

Resiliency is workload ability to recover quickly

+

and **dynamically acquire computing resources** to meet demand

Cloud Resiliency

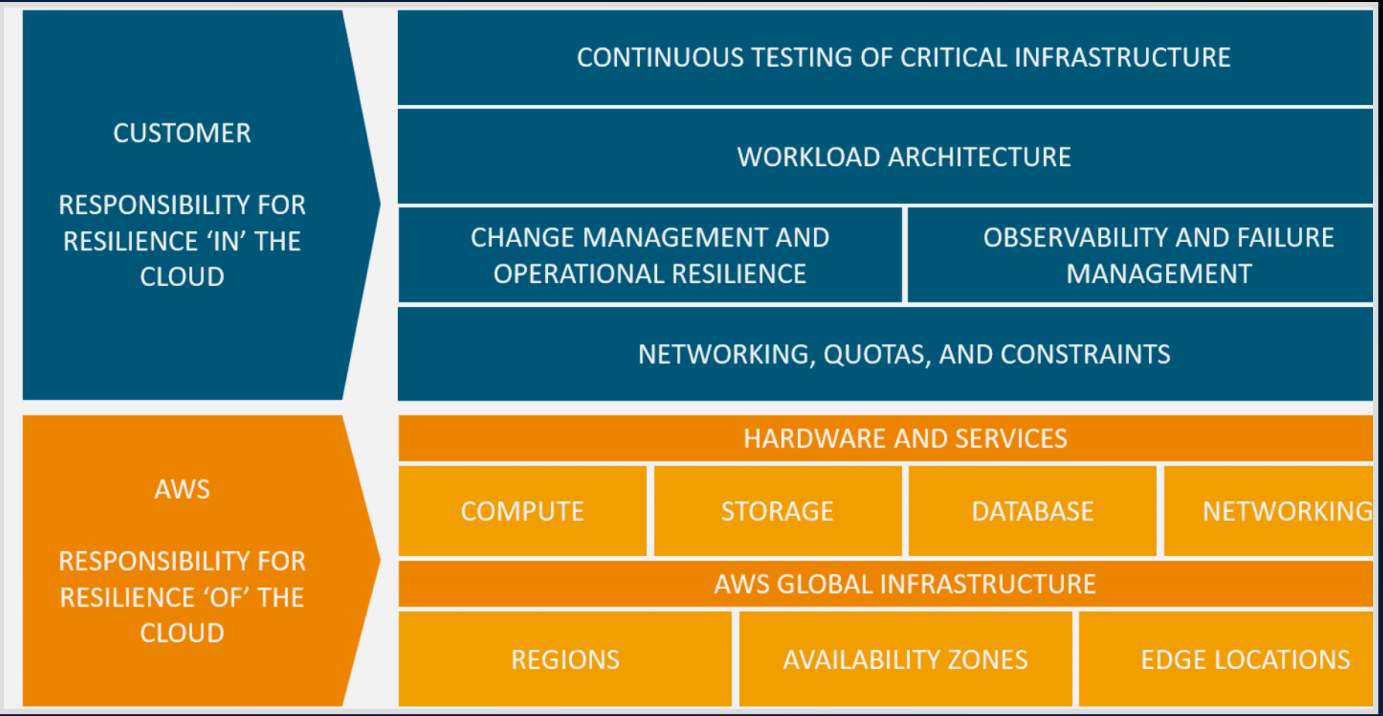

* Image: from official AWS documentation AWS original doc

Cloud Disaster Recovery (DR)

A combination of strategies and services intended to minimize downtime and data loss with fast, reliable recovery.

Cost Estimation and Cost Optimization

Cost Estimation is a process of forecasting the expenses associated with using Cloud services over a given period.

Cost Optimization is the approach to reduce expenses and maximize value while maintaining product quality and business efficiency.

Business Continuity Planning

BCP is the process involved in creating a system of prevention and recovery from potential threats to a company.

Cloud Disaster recovery is an essential part of any organization's BCP focused on safeguarding data and infrastructure

Components involved in optimizing DR costs

- Compute Costs
- Storage Costs
- Data Retrieval/Backup Costs
- Network Costs
- License Costs
- Support

Tools and Approaches for Cost Optimization

- Define Budgets and Set Alerts
- Utilize Pricing Calculators, Amazon QuickSight service
- Configure Cost Management Tools
- Continuously Review and Optimize
- 3rd party tooling, e.g.:
  - K8s -> Cast.ai
  - IaC -> Infracost

Cloud Provider: Pricing Calculators (free)

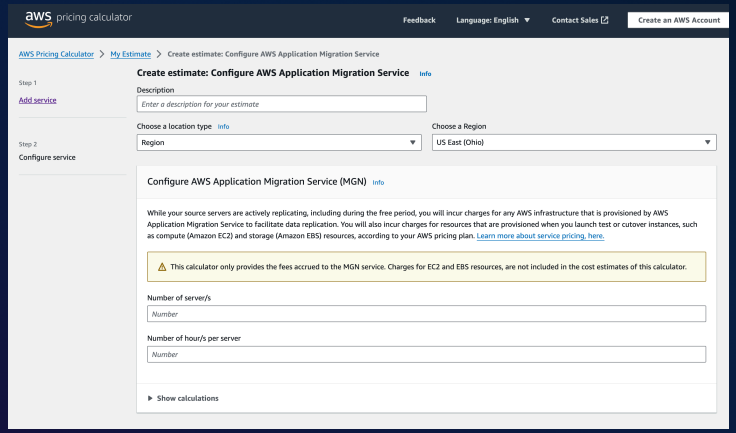

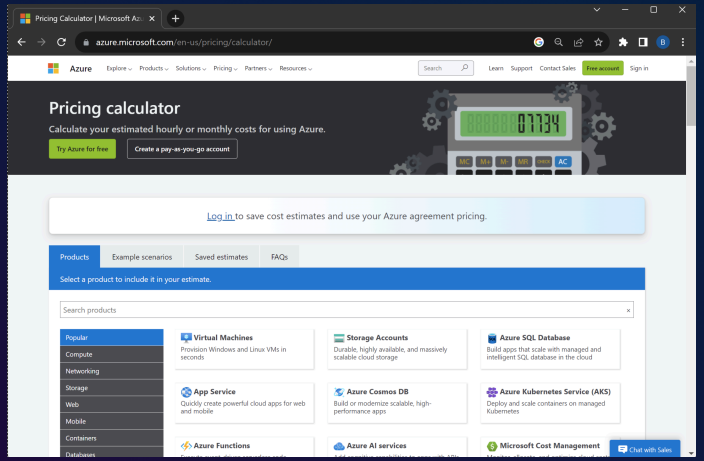

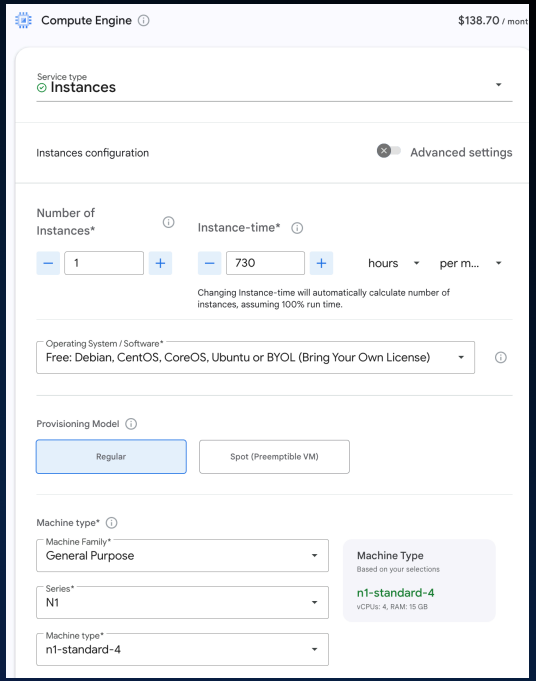

Tools and Approaches for Cost Estimation

Disaster Recovery Option	Monthly RDS Cost	Backup Storage Cost (if applicable)	Total Monthly Cost
Backup and Restore	$1,000	$100 (Backup storage)	$1,100
Pilot Light	$1,100	–	$1,100
Warm Standby	$1,500	–	$1,500
Multi-Region Failover	$2,100	–	$2,100

➜ ~ c

DR Option	Monthly EC2 Cost	Monthly EBS Cost	Total Estimated Monthly Cost
Backup and Restore	$800	$200	$1,000
Pilot Light	$900	$200	$1,100
Warm Standby	$1,200	$200	$1,400
Multi-Region Failover	$2,000	$200	$2,200

➜ ~

DR Option	Monthly VPC Cost	Monthly AZs Cost	Monthly SGs Cost	Total Estimated Monthly Cost
Backup and Restore	$100	$200	$50	$350
Pilot Light	$120	$220	$60	$400
Warm Standby	$150	$250	$70	$470
Multi-Region Failover	$200	$300	$100	$600

➜ ~

Viewing Service Utilization in Amazon QuickSight

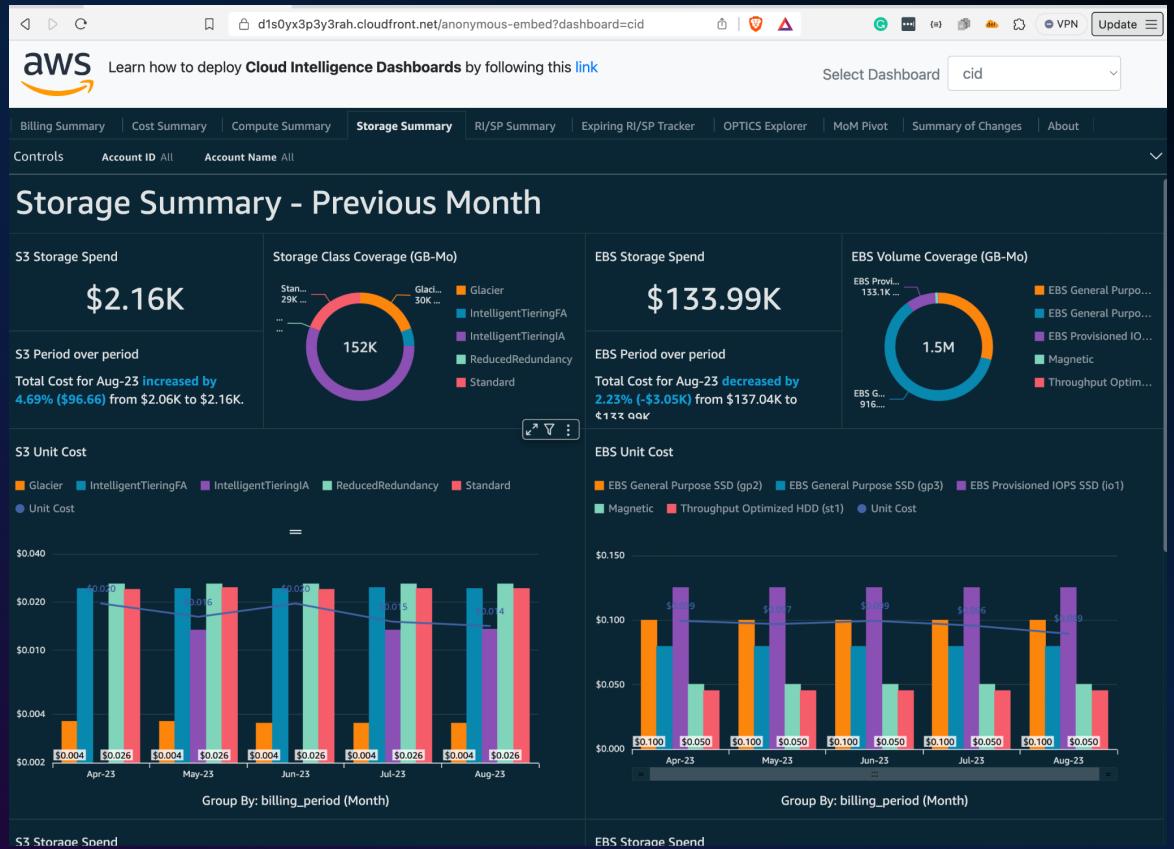

Example: Assumptions

This Example demonstrates how to optimize costs for a company operating during **US market hours**.

Goal: The disaster recovery (DR) strategy and resource utilization must ensure high availability specifically during these market hours.

Approach: Utilize *"cost factors"* to allocate monthly Cloud costs according to the resource utilization of each DR strategy.

Current DR strategy : Pilot Light (Multi Regional Failover)

Step 1: Identify and Select DR Strategies

The goal is to create a DR plan that balances cost, recovery time objectives (RTOs), and recovery point objectives (RPOs) to meet the organization's unique needs and risk profile.

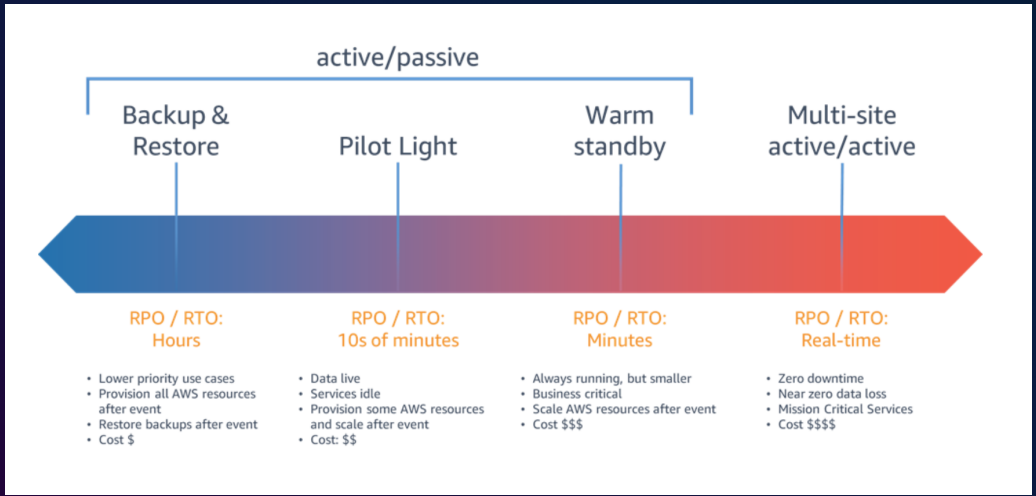

* Image: from official AWS documentation AWS original doc

Example:

Current DR strategy :

Pilot Light (Multi Regional Failover)

Is this the optimal DR strategy?

Step 2: Define DR services

- **Find Critical Workloads** - applications that are critical to business operations
- **Classify Resources by importance** - e.g., critical, important, non-essential

  Including supporting infrastructure.
- Calculate Recovery Point Objectives (**RPO**) and Recovery Time Objectives (RTO)
- Review **Compliance** and **Legal** Requirements
- Create a Recovery Hierarchy

Step 3: Calculate Utilization Percentages

- **Collect Data**:

  Collect metrics: CPU utilization for EC2 instances, read/write operations for RDS databases, etc

- **Analyze Peak and Average Utilization(***resource's potential is used)* :

  Calculate the peak and average utilization of these resources over a defined period.

- **Set Baselines:**

  Establish what is considered 100% utilization for each type of resource.

  (e.g: for EC2, it might be when the CPU is above 90% usage)

- **Calculate Percentages**:

  Utilization percentage is calculated by comparing the actual usage data to the baseline maximum capacity.

Step 4: Define “Cost Factor” for utilized resources

"Cost factors" percentage of total Cloud spending that each Disaster Recovery (DR) strategy accounts for.

$$\text{Cost Factor} = \left(\text{Total Resource Cost} / \text{Cost of DR Strategy} \right) \times 100\%$$

Example: DR infrastructure resources used

Resource	Type	Quantity	Utilization in Multi-Region Failover
EC2 Instances	c5.xlarge m5.24xlarge	2 3	Low (16%) during the day with high memory spikes High during market(trading) hours (70%), very low off-hours (5%)
EKS Clusters	m5.24xlarge	3	High during trading hours (70%), very low off-hours (5%)
RDS Databases	db.m5.large	1	Constantly high (100%)
AWS Backup	For RDS Aurora	0	0% (snapshots taken every hour)

Example#1: Cost factor for EKS resources

Objective:

- Underutilization during non-peak hours:

  Indicate over-provisioning, as usage significantly drops outside of market hours.
- High utilization during trading hours:

  Justify the choice of **m5.24xlarge** instances during those peak periods.

Example#1: Cost Factor Calculation for EKS Cluster

Total Costs:

- **Monthly cost for running**: 3 x m5.24xlarge instances, each instance costs $4.608/ per hr 24/7 30 days

  Total monthly cost = 3 x $4.608x 24 x 30 days = **$9950.4**
- **Failover-Specific Costs**: Costs associated with the Multi-Region Failover setup during peak operational hours (assuming 10 hours/day).

  Failover cost per day = 3 instances x $4.608 per hour x 10 hours = **$138.23**

  Failover cost per month = **$138.23** x 30 days = $**4146.9**
- **Cost Factor**:

  Cost Factor = (Monthly Failover Cost / Total Monthly Cost) x 100%

  Cost Factor = ($**4146.9** / $**9950.4**) x 100% ≈ **41.7%**

Approximately **41.7**% of the total EKS EC2 costs are dedicated to supporting Multi-Region Failover operations.

Is it the best DR strategy?

Example#1: Ways to Optimize EKS cost

- **Right-Sizing**: Transition to using fewer instances in standby mode during off-peak hours for EKS instances
- **Cost Efficiency**: fewer instances running continuously
- **Transition to Warm Standby**:
  - Assumption: Reduce standby instances to 3 x m5.xlarge running **24/7** in the standby region.
  - New Monthly Cost for Standby: **3 x $0.1920 x 24 x 30** = **$414.72**
  - Savings: Original Monthly Cost - New Monthly Cost = **$4146.9** - **$414.72** = **$3732.17**
  - Percentage Savings: (**$3732.17**/ **$4146.9**) x 100% ≈ 89%

Transition to a Warm Standby setup results in a reduction **89%** in monthly EC2 costs while still maintaining the necessary readiness for a quick failover during emergencies.

Example#2: Cost factor for snapshots to AWS Backup

Objective:

RDS instance with 500GB of storage, generate 30 GB of transaction logs monthly.

Backup Storage Costs:

- Storage: 500 GB x $0.095 per GB-month = **$47.50** per month.

Transaction Log Costs:

- **Assumption**: 30 GB of transaction logs generated per month.
- Storage for Logs: 30 GB x $0.095 per GB-month = $2.85 per month.

Total Monthly Cost for Continuous Backup:

- Total Cost: $47.50 (backups) + $2.85 (logs) = **$50.35** per month

Example#2: Cost Factor Calculation for AWS Backup

Cost Factor for Continuous Backup:

- Cost Factor = ($50.35/$10,000) × 100 % = 0.5035%

Operational Flexibility:

- Simplified Operations: Managing continuous backups can be simpler operationally, because it removes the need to schedule snapshot timings and manage snapshot storage growth directly.

Example#2: Ways to Optimize AWS Backup cost

- **Review and Adjust:**
  Regularly review backup storage requirements.
  Delete outdated backups and adjust retention policies.

- **Lifecycle Policies**:
  Implement lifecycle policies to transition older backups to more cost-effective storage classes, such as moving backups to Amazon S3 Glacier for long-term storage at lower costs.

- **Cross-Region Backup Optimization**: for cross-region backups, strategically select regions where storage costs are lower, if regulatory compliance permits.

Viewing Service Utilization in Amazon QuickSight

CUDOS Dashboard: EC2 Instance Usage %

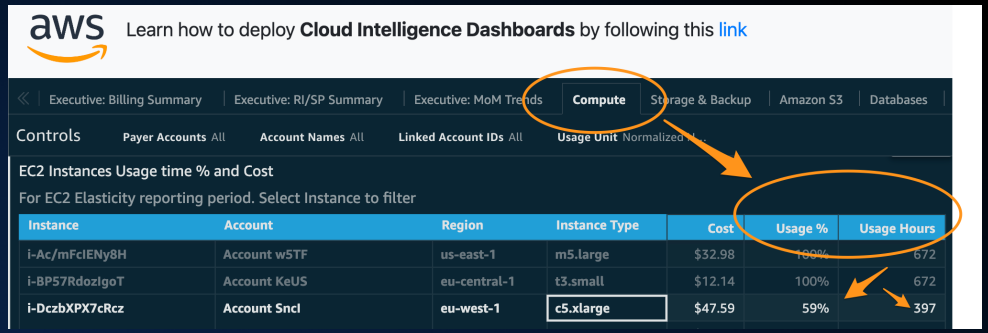

Cloud Intelligence Dashboard: Recommended Instance Family Changes

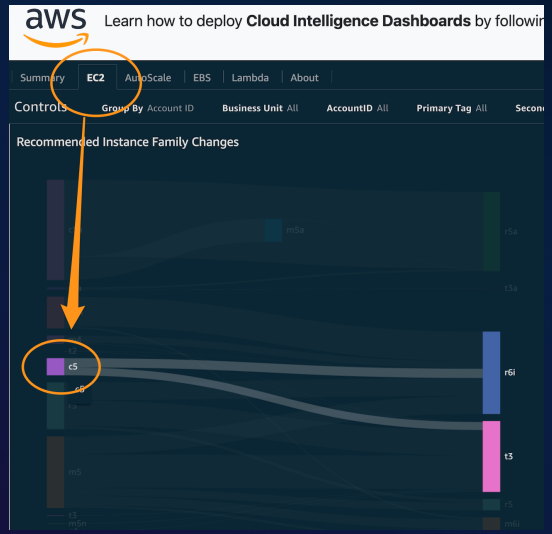

Example#3: Optimize instance type using QuickSigh

Resource	vCPUs	Memory (GiB)	Bandwidth	Pricing (on Demand)
c5.xlarge	4	8	10 Gibps	**0,1700 US$**
r6g.large	2	16	10 Gibps	0,1008 US$
t3.large	2	8	5 Gibps	$0.0832 x2 $**0.1664**

Step 1: Identify and Select DR Strategies

The goal is to create a DR plan that balances cost, recovery time objectives (RTOs), and recovery point objectives (RPOs) to meet the organization's unique needs and risk profile.

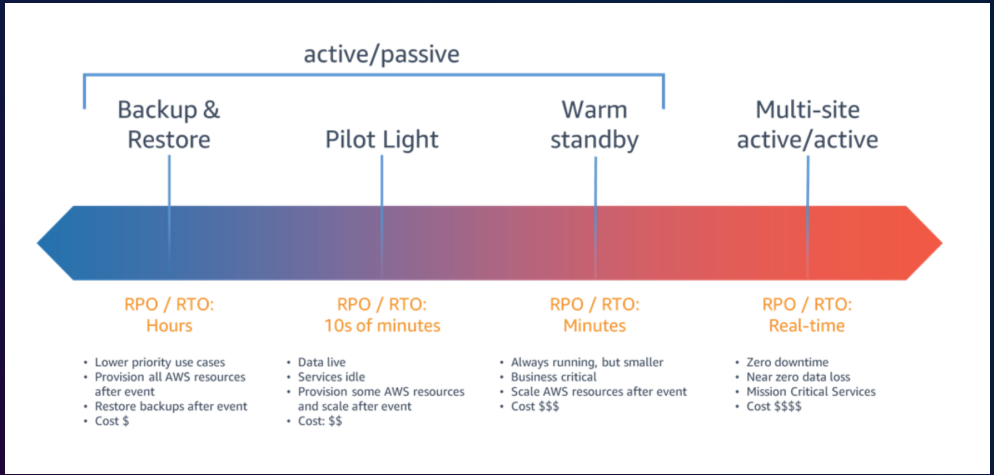

Image: from official AWS documentation AWS original doc

Current DR strategy :

Pilot Light (Multi Regional Failover)
Or
Warm Standby (Multi Regional Failover)

Choices to improve DR more

Results: DR infrastructure resources used

Resource	Type	Quantity	Utilization in Multi-Region Failover
EC2 Instances	**t3.large** **m5.xlarge**	1-2 1-3	High during the day (70%) High during trading hours (70%), very low off-hours (5%)
EKS Clusters	m5.xlarge	1-3	High during trading hours (70%), very low off-hours (5%)
RDS Databases	db.m5.large	1	Constantly high (100%)
AWS Backup	For RDS Aurora	1	0% (snapshots taken every hour)

More tasks for Cost Optimization improvements

- **Scheduled Scaling**

Implement scheduled scaling for services to automatically adjust instance counts based on predictable workload patterns. Scale down during low-demand periods to save costs.

- **Tagging and Cost Allocation**

Tagging strategies to categorize resources by DR strategy, application, or department. Cost Allocation tags help in accurate cost attribution and allocation.

- **Reserved Instances (RIs)**

Purchase Reserved Instances for predictable workloads with stable usage patterns. RIs offer significant cost savings compared to On-Demand instances.

Summary

- **Cost Analysis**: Continuously monitor and analyze disaster recovery costs to identify any unexpected or unnecessary expenses.

- **Resource Optimization**: Rightsize resources in the secondary region to ensure over-provisioning not happening

- **Automation**: Implement automation for resource provisioning and deprovisioning to minimize costs when resources are not actively in use.

THANK YOU!

Find me if any questions:

- https://in-n-out.cloud
- linkedin.com/in/login-id