



Technology, meet Humanity.

opus:interactive

# 16 Powerful Tips to Lower Your AWS Spending

**AWS is quickly becoming the default choice for public cloud providers, due in part to the many advantages it offers businesses. Chief among these are the significant cost savings. To capitalize on changing data infrastructure needs, many companies are moving away from the traditional capital-heavy data center model and replacing it with a pay-as-you-go approach. However, organizations often forget to adapt the strategy of using the cloud. Rather, they continue to operate new services with the old mindset of traditional data centers, and because of that, they lose money.**

**To optimize the cost of AWS, companies need to plan in advance and leverage the many tools that AWS provides for monitoring resource usage. By understanding how your applications are using AWS resources, you can make informed decisions about how to allocate these resources and reduce costs. In this paper, we're going to discuss the most popular, yet often overlooked, tools. This is not intended to be an ultimate guide, but it is a great place to start planning AWS usage or evaluating existing practices.**

## 1. Talk to your cloud partner

Cloud service providers always prefers businesses to use something rather than nothing. That's why they will be eager to talk, share best practices for cost optimization and provide their own instructions. Get in touch with your account representative and show them what you're up to.

In addition, small startup businesses may be entitled to AWS credits through business incubators or AWS itself. Again, talk to people, explain the action plan and ask for the credits to run testing and proof of concept (POC) implementations. On the other hand, **big companies can get a discount by reaching a certain level of resource consumption.** The key here, again, is to know an account manager and to be transparent.

### OPUS TECH TIP:

**Opus Interactive** specializes in building tailored hybrid and multi-cloud solutions for your IT ecosystem needs.

## 2. Start by using the free tier

AWS gives a small number of resources for a limited or even unlimited time. However, there are some performance, time and volume constraints. For example, EC2 instance types are just limited to t2.micro or t3.micro and DynamoDB database only comes with 25 GB of space, but it provides a great environment to test solutions before making an enterprise-wide commitment. The free tier is currently available for over 60 products. There are also three different types of offers depending on the product used:

- Always free: Doesn't expire at all
- 12 months free: Expires within 12 months after the initial signing up
- Trials: Short-term free trials

### 3. Consider regions accordingly

It's easy to notice that the price for AWS services depends on the physical location where data centers are placed. This may sound obvious but migrate resources to a region with lower prices when it makes sense.

For example, imagine your enterprise application needs to run in Europe for General Data Protection Regulation (GDPR) compliance purposes. A recent increase in post-pandemic global demand for energy, coupled with a rise in natural gas and coal prices, and a decrease in solar and wind power generation because of turbines' reduced output, caused prices for electricity to increase around the world. The trend was exacerbated with Russia's invasion of Ukraine in February 2022. **If your application was housed in a data center in Germany, you would have seen wholesale energy prices go from \$0.0349/hour in January 2020 to \$0.1775/hour in May 2022.** Whereas a data center in Sweden would have seen costs rise from \$0.0241/hour in January 2020 to \$0.0953/hour in May 2022. Comparatively, wholesale prices for the US are an average of \$0.0126/hour this past May.

Does this mean we have to abandon all highly priced regions and at once? Not necessarily. Just when it makes sense. When the company is operating in Germany and trying to provide users with the lowest latency that meets GDPR compliance, it would be impossible moving applications to the United States. Although, less demanding services or non-compliant workloads and microservices would stretch your IT budget.

#### OPUS TECH TIP:

**Opus Colocation** in the LEED Silver Hillsboro data center currently averages \$0.0942/hour with on ramps to public cloud.

## EC2 instances

### 4. Select appropriate EC2 instance types

An EC2 service would probably be one of the first picks for a cloud journey with AWS. With over 60 instance types available, picking the most suitable is often an overwhelming choice to make. Take a deep breath and think of the actual purpose of the instance. Based on this, you can narrow down the types that are the best fit by looking at the table below.

Use cases	Example	Preferred instance type
General purpose machines, which should be balanced on compute, storage and network.	Apache, NGINX, Kubernetes, Docker, VDI and development environments	T3
Compute bound applications that benefit from high performance CPU.	High performance web servers, highly scalable multiplayer gaming and video encoding	C4 and C5
Applications that process large data sets in memory.	High performance databases (e.g. SAP HANA), big data processing engines (e.g. Apache Spark or Presto) and high performance computing (HPC)	X1 and R5
Floating point number calculations, intensive graphics processing or data pattern matching.	Machine/deep learning applications, computational finance, speech recognition, autonomous vehicles or drug discovery	G4, F1 and P3
Intensive sequential read/write operations or handling large data sets.	NoSQL databases (e.g. Cassandra, MongoDB, Redis), scale-out transactional databases and data warehousing.	D2 and i3

EC2 instance type selection table.

#### OPUS TECH TIP:

Remember, right-sizing is choosing the best priced option that still meets performance requirements. A rule of thumb is to achieve 80% instance resource utilization over a long period of time.

## 5. Always prefer the most modern CPUs

Be sure to come back to the instance type table periodically, as technology doesn't freeze and CPU manufacturers are coming up with more performant and less energy-consuming CPUs almost every year and AWS implements them. For example, a simple switch from c4.xlarge to c5.xlarge for 10 instances would provide approximately \$2,500 in annual savings, while delivering more RAM (16 -> 20 GB) for each instance and around 5% better performance at the same time.

### OPUS TECH TIP:

Check the **AWS Compute Optimizer** for recommendation this type of change for optimization.

Region: US East (Ohio)						Region: US East (Ohio)					
	vCPU	ECU	Memory (GiB)	Instance Storage (GB)	Linux/UNIX Usage		vCPU	ECU	Memory (GiB)	Instance Storage (GB)	Linux/UNIX Usage
<b>Compute Optimized - Current Generation</b>						<b>Compute Optimized - Previous Generation</b>					
c5.large	2	10	4 GiB	EBS Only	\$0.085 per Hour	c4.large	2	8	3.75 GiB	EBS Only	\$0.10 per Hour
c5.xlarge	4	20	8 GiB	EBS Only	\$0.17 per Hour	c4.xlarge	4	16	7.5 GiB	EBS Only	\$0.199 per Hour
c5.2xlarge	8	39	16 GiB	EBS Only	\$0.34 per Hour	c4.2xlarge	8	31	15 GiB	EBS Only	\$0.398 per Hour
c5.4xlarge	16	73	32 GiB	EBS Only	\$0.68 per Hour	c4.4xlarge	16	62	30 GiB	EBS Only	\$0.796 per Hour
c5.9xlarge	36	139	72 GiB	EBS Only	\$1.53 per Hour	c4.8xlarge	36	132	60 GiB	EBS Only	\$1.591 per Hour

*C4 and C5 families comparison.*

## 6. Consider EC2 Instance Store for temporary data

Typically, AWS users would look at EBS storage (i.e., disks) when setting up new EC2 instances. Those disks can be attached to instances, detached from them and snapshotted for data protection cases. Whenever the instance is stopped, the data remains on the disk and doesn't go anywhere.

There is another option that is worth considering: Using local disks via the EC2 Instance Store. The major difference between these disks is that they are cleaned once the corresponding instance is stopped. They are "free" as users are paying for instance usage only.

While it's understandable that you wouldn't want to use local disks for valuable data, there are other cases when they fit perfectly, such as with temporary data like cache or logs. And don't worry, there is just enough performance, as underlying storage is as fast SSD or even NVMe. Using the EC2 Instance Store will result in less EBS consumption as well as smaller monthly bills.

## 7. Make use of spot instances

Spot instance pricing fluctuates based on the supply and demand of available unused AWS EC2 capacity. AWS users may bid their desired price. The market agrees to temporarily sell excess capacity at the lower price, resulting in a price reduction that's three to six times less than the regular price.

As the price goes above that preferred limit, provisioned instances will be automatically terminated with short notice. To avoid critical data exposure to such fluctuations, use for rendering, big data, analytics and web services behind load balancers.

## 8. Aim to predict the use of services and commit upfront

Although we talk about public cloud flexibility and the shift from CAPEX to OPEX, some of the AWS features that reduce your invoice, namely reserved instances (RIs) and saving plans, would ironically look like going back to CAPEX. However, with an actual discount of over 50%, they should be on everyone's "best practices for implementation" list.

Start with looking at RIs, which provide a discounted hourly rate and an optional capacity reservation for EC2 instances. In exchange for a commitment of one or three years, you get a discount on instance costs. On-demand instances are a good option for someone who prefers to provision workloads with unlimited flexibility. However, for constantly running workloads with a predictable load (i.e., web services), RIs are much better.

Savings summary:

- The more reservation time taken, the better the discount
- The more paid upfront, the better the discount (*but zero upfront payment is also available*)
- Standard class is cheaper than convertible (but instance type can't be changed after creation)

Watch out for convertible instances, as they can't be downsized or sold on the **AWS Marketplace**. Start with the smallest instance and upgrade when needed to avoid being left with a commitment to a monthly payment for up to 36 months.

### OPUS TECH TIP:

Be sure to check out the [AWS Pricing Calculator tool](#) (see below) when considering RIs and Saving Plans.

This instrument will be very helpful, as it shows all the potential savings and it allows for much better planning.

Comparing Savings Plans and RIs				
	Compute Savings Plans	EC2 Instance Savings Plans	Convertible RIs*	Standard RIs
Savings over On-Demand	Up to 66 percent	Up to 72 percent	Up to 66 percent	Up to 72 percent
Lower price in exchange for monetary commitment	✓	✓	—	—
Automatically applies pricing to any instance family	✓	—	—	—
Automatically applies pricing to any instance size	✓	✓	—**	—**
Automatically applies pricing to any Tenancy or OS	✓	✓	—	—
Automatically applies to Amazon ECS using Fargate	✓	—	—	—
Automatically applies to Lambda	✓	—	—	—
Automatically applies pricing across AWS Regions	✓	—	—	—
Term length options of 1 or 3 years	✓	✓	✓	✓

[RIs and Saving Plans comparison.](#)

# Object Storage

## 9. Hail to object storage

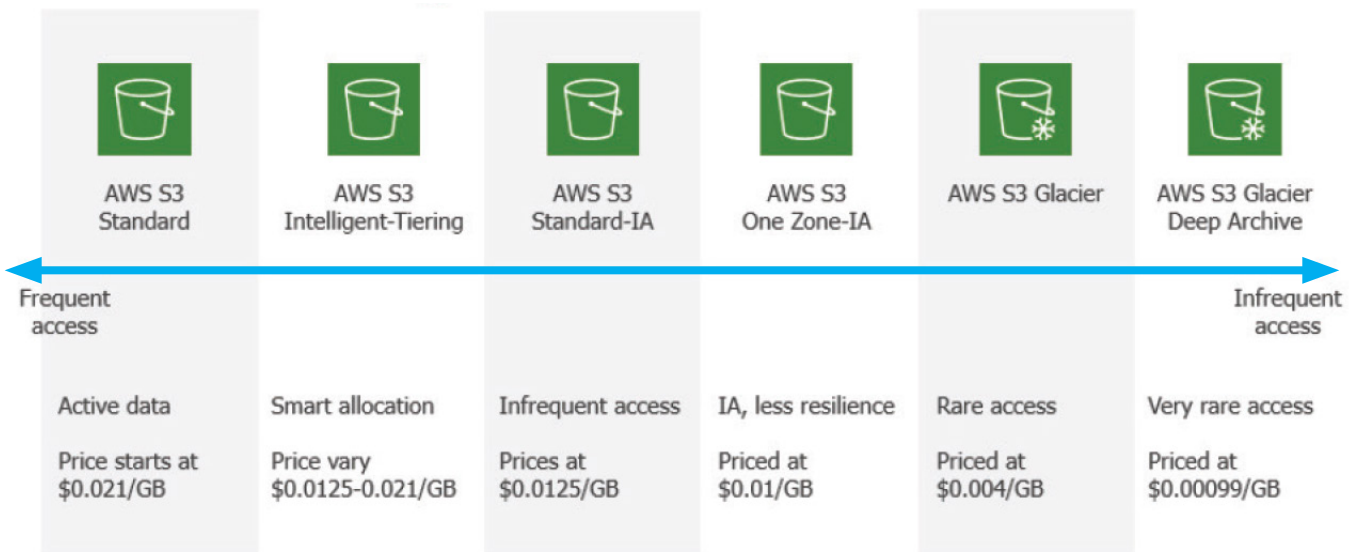
Hybrid isn't just for clouds anymore. It's storage too. AWS provides multiple storage tiers at different prices that are designed to meet requirements for performance, availability and durability. There are three broad categories of storage services offered: Object, block and file storage. Amazon's object storage offering, **Simple Storage Service (S3)**, is the most cost effective of the three storage categories. Within **Amazon S3**, you can easily move data between the storage classes further down the road to balance the frequency of access with price to optimize storage costs. All storage types have different usage scenarios and their pricing varies.

**OPUS TECH TIP:**

Data access/egress fees are an added expense in AWS and not included in storage estimate.

Full compliance, with no added egress fees, can be obtained with S3 compatible, **Opus Object Storage**. Pricing is a flat rate of \$0.025/GB.

### AWS S3 Storage classes



AWS S3 storage classes partial comparison.

The smart approach here is to combine them, depending on the type of task, object nature and access frequency. Click on desired S3 bucket, select "management", then "analytics" and then add "storage class analysis," which will be helpful to see access patterns. While it doesn't give recommendations for transitions to the One Zone-IA or S3 Glacier, it provides deeper visibility into the data.

Make sure to consider the new kid on the block: **S3 Intelligent Tiering**. For an added fee, Amazon will automatically detect access data patterns and, depending on the popularity of the object, will move them between two tiers, standard and infrequent access, all without performance impact or operational overhead. Less popular objects (i.e., ones that haven't been accessed for 30 consecutive days) will be moved to the infrequent access tier and brought back later should they be requested. Since there are no retrieval fees within this mechanism, objects can go back and forth forever. Real-world scenarios show up as having 20% savings, since the theory is partially lowered by the payment for such monitoring, yet is still too attractive to miss.

**OPUS TECH TIP:**

In AWS S3 Storage, objects that are less than 128 KB will never be transitioned to the infrequent access tier and thus will be billed at the usual rate for the frequent access tier.

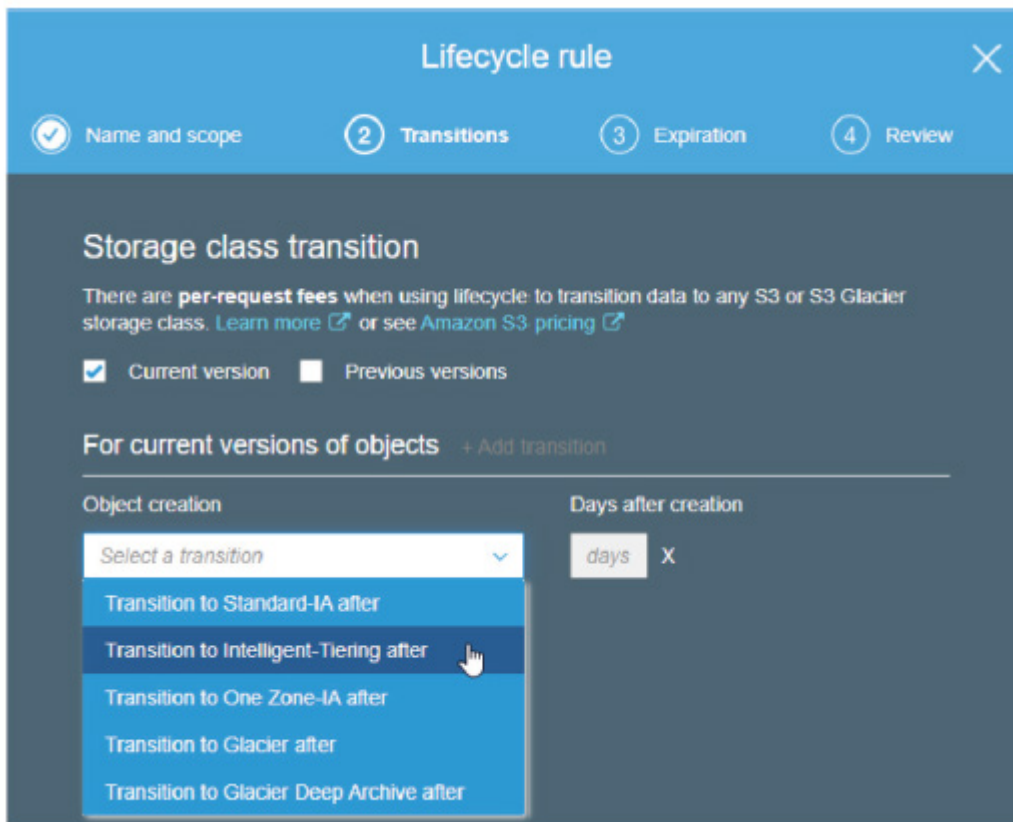
Objects that live less than 30 days, as they will be billed for a minimum of 30 days regardless.

## 10. Befriend a lifecycle policy

Take the time to learn and read the documentation on lifecycle policies. Enabling it on infrastructure can save time, cost, and resources. The policy is a combination of rules for objects that have a well-defined use pattern, so you can:

- Use lifecycle rules to manage objects throughout their lifetime
- Automate the transition to tiered storage to lower costs
- Expire objects based on retention needs or service-level agreements (SLAs)

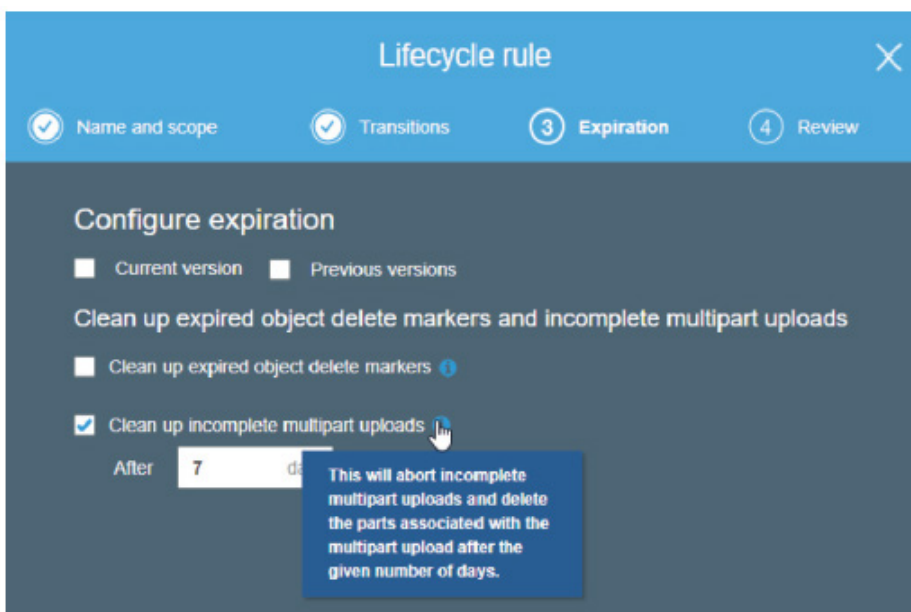
This is a very powerful tool when configured properly. Learn the difference between the S3 storage classes and start playing with lifecycle rules to better fit organization needs.



*Lifecycle rule creation.*

## 11. Clean up multipart incomplete uploads

S3's multipart upload feature is enabled by default and accelerates the upload of large objects by splitting them up into logical parts that can be uploaded in parallel. The issue comes in when those uploads never finish for some reason. The actual incomplete data won't be visible in the bucket, nor will it be automatically deleted, so you won't notice a thing, except when the cost incurs in larger monthly bills. To prevent this, go to the "bucket management" settings, create a new lifecycle rule and enable the "clean up incomplete multipart uploads" option.



S3 bucket. Expiration configuration.

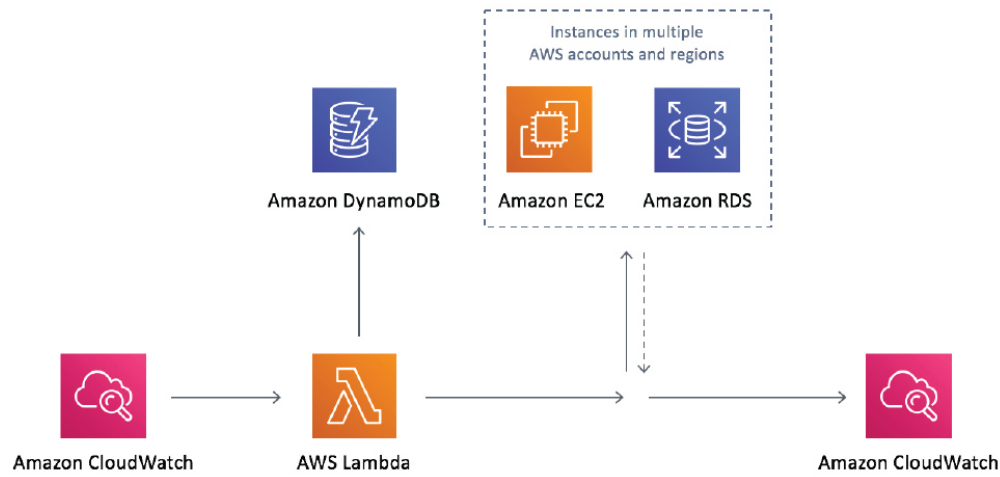
## Monitoring and management

Hybrid and multi-cloud strategies are complex. AWS has developed some tools, each targeting a certain aspect of price reduction to help companies on their cloud spend optimization journey.

## 12. Automate for unused

Idle resources can be a major cost contributor to your AWS bill. Letting unused instances and databases sit idle means accruing charges for something that isn't used. For example, if you have a development environment that's only accessed during the day on weekdays, the best option would be not to run it 24/7. By programmatically stopping it at night and starting it again the following morning and at the end of the week, the costs could be cut almost in half. It's also good strategy to enable **Amazon CloudWatch** alarms to automatically stop or terminate instances that have been idle for longer than a specified period.





Amazon CloudWatch diagram.

## 13. Bring the AWS Pricing Calculator into action

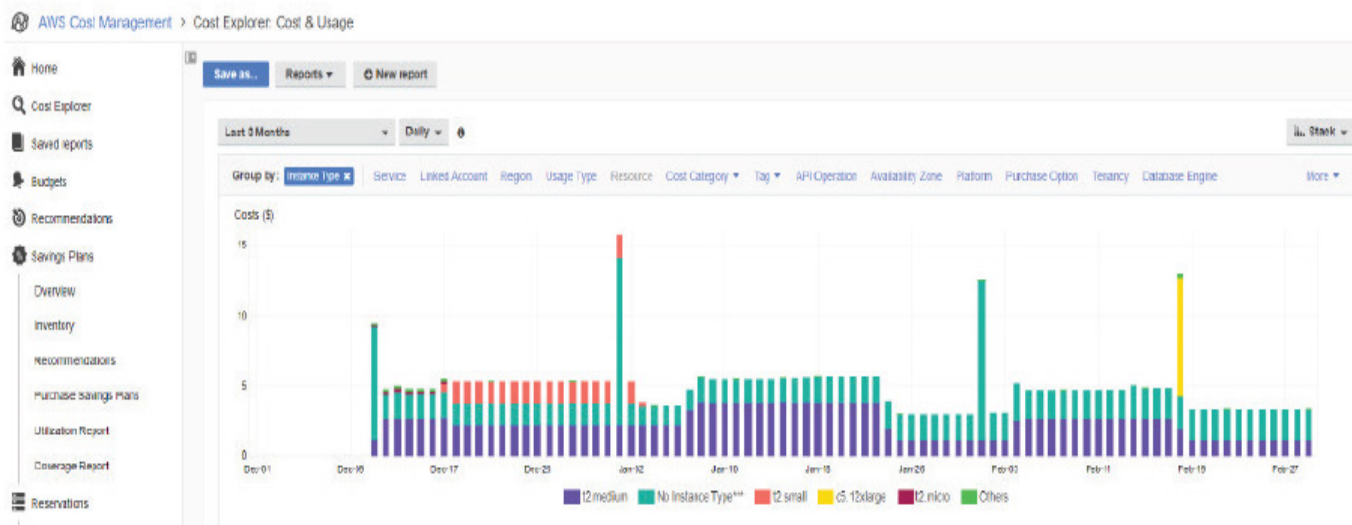
The **AWS Pricing Calculator** allows you to estimate account spending for over 40 AWS services. It can be very helpful when deciding between other competitive strategies as well. For example, when switching from on-demand EC2 instances to RIs and Savings Plan offerings, you can check the discount each feature provides. With every value added to the table, it will show a resulted calculation and help to select the actual service.

The screenshot shows the AWS Pricing Calculator interface. At the top, there is a navigation bar with the AWS logo, 'pricing calculator', 'Feedback', 'English', and 'Cor'. Below this is a 'Pricing strategy Info' section. It contains three columns of options: 'Pricing model' with radio buttons for EC2 Instance Savings Plans (selected), Compute Savings Plans, Standard Reserved Instances, Convertible Reserved Instances, and On-Demand Instances; 'Reservation term' with radio buttons for 1 Year (selected) and 3 Year; and 'Payment options' with radio buttons for No Upfront (selected), Partial Upfront, and Full Upfront. Below these options is a 'Show calculations' section. It includes 'Unit conversions' for EC2 Instance Savings Plans, showing a rate of 0.0118 USD per hour for a 1-year term with no upfront cost. It calculates a total commitment of 103,368 USD and an hourly cost of 0.0118 USD. Finally, it shows 'Pricing calculations' for 100 instances over 730 hours in a month, resulting in a monthly instance savings cost of 861.40 USD.

Amazon CloudWatch diagram.

## 14. Monitor consumption with standard tools

**AWS Cost Explorer** should be your go-to, since this instrument makes it very easy to visualize where all the money goes and address the services that consume the most. This can be found based on the analysis conducted, identifying interesting patterns and drilling down to the root cause.



Cost Explorer. View instance types consumption.

**AWS Compute Optimizer** provides an overview of optimization opportunities for AWS resources based on the data that's been collected and analyzed for the account (or all the accounts under the master one).

The screenshot shows the AWS Compute Optimizer interface. The main content area is titled 'Recommendations for EC2 instances (1)'. Below the title, there are filters for 'Regions: US East (N. Virginia)' and 'All findings'. A table lists the recommendations:

Instance ID	Instance name	Finding	Current instance type	Current On-Demand price	Recommended instance type	Recommended On-Demand price	Price difference
i-0bc3f7236476d4b33	VMA BYO	Optimized	t2.medium	\$0.084 per hour	t2.medium	\$0.016 per hour	-\$0.068 per hour

AWS Compute Optimizer. Recommendations for EC2 instances.

## 15. Consolidate users with master account

If multiple people in the enterprise need to operate AWS, consolidated billing will not only make payments easy but can also reach a certain threshold that once reached, provides a discount on consumed resources like S3. Some AWS tiers also reward higher usage with lower prices and discounted rates for purchasing instances in advance (like the RIs and Saving Plans from above). In addition, unused resources can be redistributed from one child account to another. Apply the previously mentioned cost optimization tools here and prevent organization chaos when operating numerous accounts.

## 16. Don't forget about networking costs

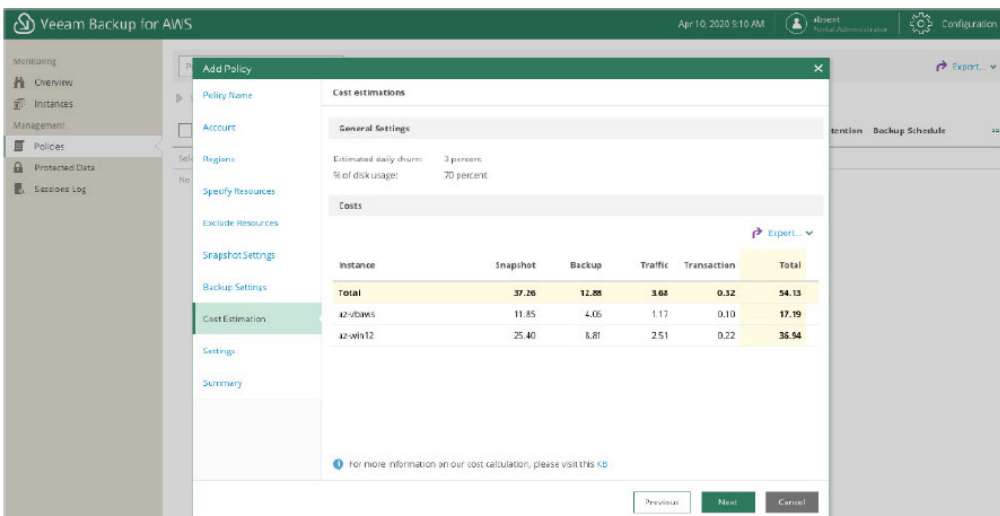
- In case of massive traffic exchange between on-premises sites and AWS, the Direct Connection feature will help you get more consistent network experience, increase bandwidth throughput and secure connectivity.
- Static content (i.e., images, videos, music) is better and more cheaply distributed through S3 and the CloudFront combination. With a great set of edge servers in different areas of the world, end users will get the data coming from the service cache that's located closer to them.
- Analyze the traffic flow between different AWS services. Configure VPC endpoints so the traffic from VPC to other services like S3 would go through the endpoint, which bypasses the internet and public networks and results in a more secure and cheaper connection.
- Don't forget about inter-availability zone traffic, as that will be another expense. Reconsider fault-tolerance architecture, it might not be needed for all services, or can be achieved through other technologies.

### OPUS TECH TIP:

**Opus AWS Management Services** includes full management of your AWS ecosystem. Additional colocation, tailored IaaS, single-pane hybrid and multicloud monitoring, and S3 compatible object storage are also available for end-to-end IT.

## Bonus tech tip

**Opus Veeam Backup for AWS** can take backups of EC2 instances to ensure its protection and perform additional operations, such as restoring systems back to AWS or even on-premises, to move data around the infrastructure as required. Moreover, it will always show a predicted AWS cost amount caused by these operations, so any protection strategy can be adjusted accordingly.

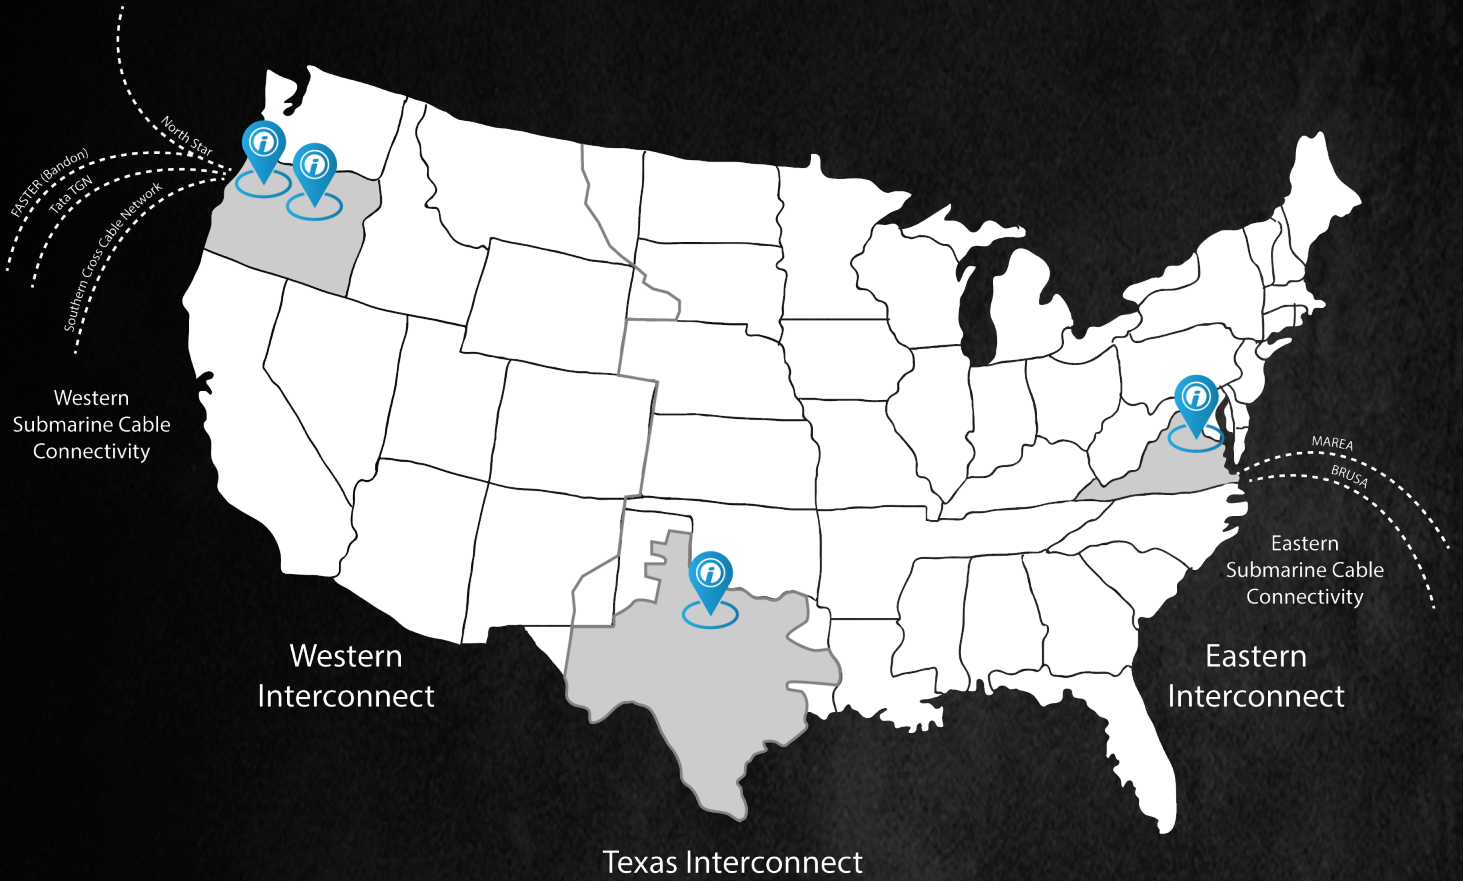


Instance	Snapshot	Backup	Traffic	Transaction	Total
<b>Total</b>	<b>37.26</b>	<b>12.88</b>	<b>3.68</b>	<b>0.32</b>	<b>54.13</b>
i2-vbaws	11.85	4.06	1.17	0.10	17.19
i2-wh12	25.40	8.81	2.51	0.22	36.94

Veeam Backup for AWS.

### OPUS TECH TIP:

**Opus Veeam Backup for AWS** offering is S3 compatible. Additionally, the solutions are tailored to your RPO and RTO needs at world class data centers strategically located in separate regions and separate energy grids.




## Opus Interactive


Oregon: 8135 NE Evergreen Parkway Hillsboro, OR 97124  
 Northern Virginia: 11680 Hayden Road, Manassas, VA 20109

☎ (503) 972-6690 main  
 🌐 [opusinteractive.com](http://opusinteractive.com)

Oregon  
 Datacenter



NoVA  
 Datacenter



Founded in 1996, Opus Interactive has earned a reputation for custom hybrid cloud solutions that fit unique requirements for security, scalability, cost, and future growth needs of its customers. An accredited member of the International Managed Services Provider Alliance, the Company operates from Tier III+ data centers located in Hillsboro, Portland, Dallas, and Northern Virginia. Through close partnerships with industry-leaders and a commitment to customer satisfaction, Opus offers Hybrid Cloud & IaaS (with IaaS, Object Storage, and Virtual Desktop Infrastructure), Enterprise Colocation, IT Services, Network Connectivity, DRaaS & Backup, Hybrid Cloud Monitoring & Management.

OPUS Interactive is a woman and minority-owned enterprise that has worked closely with VMware and HPE partnership programs since 2005. With past performance that includes more than 20 years of proven results and current compliance with PCI-DSS, HIPAA, and SSAE 18 SOC 2, Opus helps customers reduce cost and optimize resources using efficient operations.