



## How AWS Pricing Works

*July 2014*

*(Please consult <http://aws.amazon.com/whitepapers/> for the latest version of this paper)*

## Table of Contents

Table of Contents .....	2
Abstract .....	3
Introduction.....	3
Fundamental Pricing Characteristics .....	4
Amazon Elastic Compute Cloud (Amazon EC2) .....	5
Amazon Simple Storage Service (Amazon S3) .....	6
Amazon Elastic Block Store (Amazon EBS).....	6
Amazon Relational Database Service (Amazon RDS).....	7
Amazon CloudFront .....	7
How to Further Save Costs .....	8
Conclusion .....	10
Resources .....	11
Cost Calculation Example .....	12

## Abstract

While the number and types of services offered by AWS has increased dramatically, our philosophy on pricing has not changed: at the end of each month, you pay only for what you use, and you can start or stop using a product at any time. No long-term contracts are required.

Pricing information for each service on our website is available at <http://aws.amazon.com/pricing/>. Our strategy of pricing each service independently gives you tremendous flexibility to choose the services you need for each project and to pay only for what you use.

This whitepaper will help you understand how to effectively estimate the costs of running your specific project on AWS. We provide several examples that leverage the [AWS Simple Monthly Calculator](#). For each example, this paper discusses its architecture, example usage of each service, cost breakdown for each service, and the total estimated monthly charge.

## Introduction

AWS offers a range of cloud computing services. For each service, you pay for exactly the amount of resources you actually need. This utility-style pricing model is explained below:

- **Pay as you go.** No minimum commitments or long-term contracts required. You replace your upfront capital expense with low variable cost and pay only for what you use. There is no need to pay upfront for excess capacity or get penalized for under-planning. For compute resources, you pay on an hourly basis from the time you launch a resource until the time you terminate it. For data storage and transfer, you pay on a per gigabyte basis. We charge based on the underlying infrastructure and services that you consume. You can turn off your cloud resources and stop paying for them when you don't need them.
- **Pay less when you reserve.** For certain products, you can invest in reserved capacity. In that case, you pay a low upfront fee and get a significantly discounted hourly rate, which results in overall savings up to 60% (depending on the type of instance you reserve) over equivalent On-Demand capacity.
- **Pay even less per unit by using more.** You save more as you grow bigger. For storage and data transfer, pricing is tiered. The more you use, the less you pay per gigabyte. For compute, you get volume discounts up to 10% when you reserve more.
- **Pay even less as AWS grows.** Most importantly, we are constantly focused on reducing our data center hardware costs, improving our operational efficiencies, lowering our power consumption, and generally lowering the cost of doing business. These optimizations and AWS's substantial and growing economies of scale

### AWS Free Tier

To help new AWS customers get started in the cloud, AWS offers a free tier. If you're a new AWS customer, you can run a free Amazon EC2 Micro Instance for a year while also leveraging the free tier for select portions of Amazon S3, Amazon Elastic Block Store, Amazon Elastic Load Balancing, AWS data transfer and other AWS services. For more information, go to <http://aws.amazon.com/free>

result in passing savings back to you in the form of lower pricing. Since 2006, AWS has consistently lowered prices (45 price drops as of Aug. 1, 2014). To see the latest pricing visit the [AWS pricing page](#)

- **Custom pricing.** What if none of our pricing models work for your project? Custom pricing is available for high volume projects with unique requirements. For assistance, [contact us](#) to speak with a sales representative.

AWS also offers a variety of services for **no additional charge**:

- [Amazon VPC](#): Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the Amazon Web Services (AWS) Cloud where you can launch AWS resources in a virtual network that you define.
- [AWS Elastic Beanstalk](#): AWS Elastic Beanstalk is an even easier way for you to quickly deploy and manage applications in the AWS cloud.
- [AWS CloudFormation](#): AWS CloudFormation gives developers and systems administrators an easy way to create a collection of related AWS resources and provision them in an orderly and predictable fashion.
- [AWS Identity and Access Management \(IAM\)](#): AWS IAM controls your users' access to AWS services and resources.
- [Auto Scaling](#): Auto Scaling automatically adds or removes [Amazon EC2](#) instances according to conditions you define. With Auto Scaling, the number of Amazon EC2 instances you're using increases seamlessly during demand spikes to maintain performance, and decreases automatically during demand lulls to minimize costs.
- [AWS OpsWorks](#): AWS OpsWorks is an application management service that makes it easy to deploy and operate applications of all shapes and sizes.
- Moreover, you can consolidate all your accounts using [Consolidated Billing](#) and get tiering benefits.

## Fundamental Pricing Characteristics

There are three fundamental characteristics you pay for with AWS: compute, storage, and data transfer out. These characteristics vary slightly depending on the AWS product you are using. However, fundamentally these are the core characteristics that have the greatest impact on cost.

Although you are charged for data transfer out, there is no charge for inbound data transfer or for data transfer between other Amazon Web Services within the same region. The outbound data transfer is aggregated across AWS services and then charged at the outbound data transfer rate. This charge appears on the monthly statement as AWS Data Transfer Out.

The rest of this section breaks down the pricing characteristics for four commonly used AWS products: Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (S3), Amazon Elastic Block Store (EBS), Amazon Relational Database Service (Amazon RDS), and Amazon CloudFront.

Pricing is available for each AWS product and its specific pricing characteristics at <https://aws.amazon.com/pricing/>.

### Inbound Data Transfer at No Additional Cost

There is no charge for inbound data transfer across all Amazon Web Services in all regions. There are minimal outbound data transfer charges between Amazon Web Services within the same region.

## Amazon Elastic Compute Cloud (Amazon EC2)

---

[Amazon Elastic Compute Cloud \(Amazon EC2\)](#) is a web service that provides resizable compute capacity in the cloud. Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides complete control of your computing resources on Amazon's proven computing environment. Amazon EC2 changes the economics of computing by charging you only for capacity that you actually use.

When you begin to estimate the cost of using Amazon EC2, you need to consider the following:

- **Clock Hours of Server Time** - Resources incur charges when they are running. For example, from the time Amazon EC2 instances are launched until they are terminated, or from the time Elastic IPs are allocated until the time they are de-allocated.
- **Machine Configuration** - Consider the physical capacity of the Amazon EC2 instance you choose. Instance pricing varies with the AWS region, OS, number of cores, and memory.
- **Machine Purchase Type** - With On-Demand Instances, you pay for compute capacity by the hour with no required minimum commitments. Reserved Instances give you the option to make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly usage charge for that instance. With Spot Instances, you can bid for unused Amazon EC2 capacity. For more information how to further save costs using Reserved Instances and Spot Instances, see [How to Further Save Costs](#).
- **Number of Instances** - You can provision multiple instances of your Amazon EC2 and Amazon EBS resources to handle peak loads.
- **Load Balancing** - An Elastic Load Balancer can be used to distribute traffic among Amazon EC2 instances. The number of hours the Elastic Load Balancer runs and the amount of data it processes contribute to the monthly cost.
- **Detailed Monitoring** - You can use [Amazon CloudWatch](#) to monitor your EC2 instances. By default, basic monitoring is enabled (and available at no additional cost); however, for a fixed monthly rate, you can opt for detailed monitoring, which includes seven preselected metrics recorded once a minute. Partial months are charged on an hourly pro rata basis, at a per instance-hour rate.
- **Auto Scaling** - Auto Scaling automatically adjusts the number of Amazon EC2 instances in your deployment according to conditions you define. This service is available at no additional charge beyond Amazon CloudWatch fees.
- **Elastic IP Addresses** - You can have one Elastic IP (EIP) address associated with a running instance at no charge.
- **Operating Systems and Software Packages** - Operating System prices are included in the instance prices. To view a list of currently available operating systems that you can use with your Amazon EC2 instances and their prices, go to [Amazon Elastic Compute Cloud \(Amazon EC2\)](#). There are no additional licensing costs to run the following commercial operating systems: [Red Hat Enterprise Linux](#), [SUSE Enterprise Linux](#), [Windows Server](#), and [Oracle Enterprise Linux](#). Additionally, AWS has made it easy for you and has partnered with Microsoft, IBM and several other vendors so you can run commercial software packages on your Amazon EC2 instances. For example, [Microsoft SQL Server on Windows](#), [IBM Software](#). For commercial software packages that AWS does not provide, such as nonstandard operating systems, Oracle Applications, Windows Server applications such as Microsoft SharePoint and Microsoft Exchange, you need to obtain a license from the vendors. You can also bring your existing license to the cloud through specific vendor programs such as [Microsoft License Mobility through Software Assurance Program](#).

For more information about Amazon EC2 pricing, go to <http://aws.amazon.com/ec2/>.

## Amazon Simple Storage Service (Amazon S3)

---

[Amazon Simple Storage Service \(Amazon S3\)](#) is storage for the Internet. It provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.

When you begin to estimate the cost of Amazon S3, you need to consider the following:

- **Storage Class** - Standard Storage is designed to provide 99.999999999% durability. Reduced Redundancy Storage (RRS) is a storage option within Amazon S3 that you can use to reduce your costs by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3's standard storage. Reduced Redundancy Storage is designed to provide 99.99% durability. Each class has different rates.
- **Storage** - The number and size of objects stored in your Amazon S3 buckets as well as type of storage.
- **Requests** - The number and type of requests. GET requests incur charges at different rates than other requests, such as PUT and COPY requests.
- **Data Transfer** - The amount of data transferred out of the Amazon S3 region.

For more information about Amazon S3 pricing, go to <http://aws.amazon.com/pricing/s3/>.

## Amazon Elastic Block Store (Amazon EBS)

---

Amazon Elastic Block Store (EBS) provides block level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance. They are analogous to virtual disks in the cloud. Amazon EBS provides three volume types: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. The three volume types differ in performance characteristics and cost, so you can choose the right storage performance and price for the needs of your applications.

- **Volume Size** – Volume storage for all EBS volume types is charged by the amount you provision in GB per month, until you release the storage.
- **Volume Type** – You can choose what type of storage volume you would like. The default is a general purpose SSD. You can choose a magnetic (spinning disk) volume as a lower cost option for lower I/O needs, or Provisioned IOPS for high performance and specific I/O needs.
- **Input Output Operations per Second (IOPS)** - I/O is included in the price of General Purpose (SSD) volumes, while for EBS Magnetic volumes, I/O is charged by the number of requests you make to your volume. With Provisioned IOPS (SSD) volumes, you are also charged by the amount you provision in IOPS (multiplied by the percentage of days you provision for the month).
- **Snapshot** - Amazon EBS provides the ability to back up snapshots of your data to Amazon S3 for durable recovery. If you opt for EBS snapshots, the added cost is per GB-month of data stored.

For more information about Amazon Elastic Block Store (EBS) pricing, go to <http://aws.amazon.com/pricing/EBS/>.

## Amazon Relational Database Service (Amazon RDS)

---

[Amazon Relational Database Service \(Amazon RDS\)](#) is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, so you can focus on your applications and business.

When you begin to estimate the cost of Amazon RDS, you need to consider the following:

- **Clock Hours of Server Time** - Resources incur charges when they are running. For example, from the time you launch a DB instances until you terminated the DB instance.
- **Database Characteristics** - The physical capacity of the database you choose will affect how much you are charged. Database characteristics vary depending on the database engine, size, and memory class.
- **Database Purchase Type** - When you use On-Demand DB Instances, you pay for compute capacity for each hour your DB Instance runs, with no required minimum commitments. With Reserved DB Instances, you can make a low, one-time, up-front payment for each DB Instance you wish to reserve for a 1-year or 3-year term.
- **Number of Database Instances** - With Amazon RDS, you can provision multiple DB instances to handle peak loads.
- **Provisioned Storage** - There is no additional charge for backup storage of up to 100% of your provisioned database storage for an active DB Instance. After the DB Instance is terminated, backup storage is billed per gigabyte per month.
- **Additional Storage** - The amount of backup storage in addition to the provisioned storage amount is billed per gigabyte per month.
- **Requests** - The number of input and output requests to the database.
- **Deployment Type** - You can deploy your DB instance to a single Availability Zone (analogous to a stand-alone data center) or multiple Availability Zones (analogous to secondary data center for enhanced availability and durability). Storage and I/O charges vary, depending on the number of Availability Zones you deploy to.
- **Data Transfer** - Inbound data transfer is free, and outbound data transfer costs are tiered.

Depending on the needs for your application, it's possible to optimize your costs for Amazon RDS database instances by purchasing reserved Amazon RDS database instances. To purchase Reserved Instances, you make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly usage charge for that instance.

For more information about Amazon RDS pricing, go to <http://aws.amazon.com/pricing/rds/>.

## Amazon CloudFront

---

[Amazon CloudFront](#) is a web service for content delivery. It integrates with other Amazon Web Services to give you an easy way to distribute content to end users with low latency, high data transfer speeds, and no required minimum commitments.

When you begin to estimate the cost of Amazon CloudFront, you need to consider the following:

- **Traffic Distribution** - Data transfer and request pricing vary across geographic regions, and pricing is based on the edge location through which your content is served.
- **Requests** - The number and type of requests (HTTP or HTTPS) made and the geographic region in which the requests are made.
- **Data Transfer Out** - The amount of data transferred out of your Amazon CloudFront edge locations.

For more information about Amazon CloudFront pricing, go to <http://aws.amazon.com/pricing/cloudfront/>.

## How to Further Save Costs

Depending on the needs of your application, you can optimize your costs for Amazon EC2 instances by purchasing EC2 Reserved Instances or Spot Instances. On-Demand Instances are a good option if you run your Amazon EC2 Instances a couple of hours a day or a few days per week; however, if you plan to run your Amazon EC2 Instances more than that, Reserved Instances can save you money. To obtain Reserved Instances, you make a low, one-time payment for each instance you want to reserve, and in turn you receive a significant discount on the hourly usage charge for that instance.

The following tables show an example of potential savings for Amazon EC2 Reserved Instances for a 1-year and 3-year term running a single m3.large Linux instances in the US East Region. The prices are taken as of the date of publication and may change.

**Table 1: Savings Comparison of 1 Year Reserved Instances over On-Demand Instances**

Utilization Rate	On-Demand	1 Year Medium Utilization	1 Year Heavy Utilization
10%	\$122.98	-234%	-525%
20%	\$245.95	-86%	-212%
30%	\$368.93	-37%	-108%
40%	\$491.90	-13%	-56%
50%	\$614.88	2%	-25%
60%	\$737.86	12%	-4%
70%	\$860.83	19%	11%
80%	\$983.81	24%	22%
90%	\$1,106.78	28%	31%
100%	\$1,229.76	31%	38%

*Utilization Rate = % of time your instance is running; Prices shown for US East Region as of July 20<sup>th</sup> 2014*

As shown in table 1, if your Amazon EC2 instance is running at more than 50% utilization or more than 6 months over a 1 year term, you can achieve savings over On-Demand instances. Using Heavy Utilization Reserved Instances, you can save up to 38% over a 1 year term.



**Table 2: Savings Comparison of 3 Year Reserved Instances over On-Demand Instances**

Utilization Rate	On-Demand	3 Yr. Medium Utilization	3 Yr. Heavy Utilization
10%	\$368.93	-85%	-297%
20%	\$737.86	-8%	-98%
30%	\$1,106.78	18%	-32%
40%	\$1,475.71	31%	1%
50%	\$1,844.64	38%	21%
60%	\$2,213.57	43%	34%
70%	\$2,582.50	47%	43%
80%	\$2,951.42	50%	50%
90%	\$3,320.35	52%	56%
100%	\$3,689.28	54%	60%

*Utilization Rate = % of time your instance is running; Prices shown for US East Region as of July 20<sup>th</sup> 2014*

As shown in table 2, if your Amazon EC2 instance is running at more than 30% utilization or more than 10.8 months over a 3 year term, you can achieve savings over On-Demand instances. Using Heavy Utilization Reserved Instances, you can save up to 60% over a 3 year term.

Suppose you have deployed your web application on AWS, and you need to run six Amazon EC2 instances:

- Two m3.medium Linux Instances for the web servers
- Two m3.medium Linux Instances for the application servers
- Two m3.large Linux Instances for the database servers

You have a number of options for running these instances. For example, you can run On-Demand Instances only, On-Demand and Reserved Instances (Medium Utilization), or if you know the utilization of your Amazon EC2 instances, you can run On-Demand and a mixture of different Utilization Reserved Instances. There are a number of other options, but this table shows four different options using these three combinations.

Description	Option 1	Option 2	Option 3	Option 4
Web servers	2 On-Demand	2 On-Demand	1 On-Demand and 1 Medium Utilization RI	2 Heavy Utilization RI
App servers	2 On-Demand	2 On-Demand	1 On-Demand and 1 Medium Utilization RI	2 Heavy Utilization RI
Database servers	2 On-Demand	2 Medium Utilization RI	2 Medium Utilization RI	2 Heavy Utilization RI

The following table compares costs for each option and shows the total savings over On-Demand instances for 1-year and 3-year terms. You save money when you use Medium Utilization Reserved Instances and then you save even more money when you use Heavy Utilization Reserved Instances. Heavy Utilization Reserved Instances offer you the maximum savings over On-Demand instances. The prices are taken as of the date of publication and may change.

		Option 1	Option 2	Option 3	Option 4
		<a href="#">Calculator</a>	<a href="#">Calculator</a>	<a href="#">Calculator</a>	<a href="#">Calculator</a>
Monthly Cost	1 Year Term	\$409.92	\$285.48	\$222.54	\$106.90
	3 Year Term	\$409.92	\$267.92	\$197.66	\$87.84
One-Time Cost (Reserved Instances)	1 Year Term	-	\$724.95	\$1086.00	\$1774.00
	3 Year Term	-	\$1142.00	\$1714.00	\$2694.00
Total Cost	1 Year Term (x12)	\$4919.04	\$4150.71	\$3756.48	\$3056.80
	3 Years Term (x36)	\$14757.12	\$11857.76	\$8829.76	\$5856.24
Savings (Over Option 1)	<b>1 Year Term</b>	<b>n/a</b>	<b>15%</b>	<b>24%</b>	<b>37%</b>
	<b>3 Years Term</b>	<b>n/a</b>	<b>20%</b>	<b>40%</b>	<b>60%</b>

Prices shown for US East Region as of July 20<sup>th</sup> 2014

Your Reserved Instance will be available for the operating system (Linux/UNIX or Windows) and Availability Zone in which you purchased it. For more information about Reserved Instances, go to [Amazon EC2 Reserved Instance](#) pages.

Spot Instances are unused Amazon EC2 capacity that you bid for. Instances are charged at Spot Price, which is set by Amazon EC2 and fluctuates periodically depending on the supply of, and demand for, Spot Instance capacity. If your maximum bid exceeds the current Spot Price, your bid request is fulfilled, and your instances will run until either you choose to terminate them or the Spot Price increases above your maximum bid, whichever is sooner. To learn more about Spot Instances, go to <http://aws.amazon.com/ec2/spot-instances/>.

## Conclusion

While the number and types of services offered by AWS has increased dramatically, our philosophy on pricing has not changed. You pay as you go, pay for what you use, pay less as you use more, and pay even less when you reserve capacity. Projecting costs for a use case, such as web application hosting, can be challenging, because a solution typically uses multiple features across multiple AWS products, which in turn means there are more factors and purchase options to consider.

The best way to estimate costs is to examine the fundamental characteristics for each AWS product, estimate your usage for each characteristic, and then map that usage to the prices posted on the website. To help you understand how AWS pricing works in the context of real-world solutions, see a worked out example later in this paper.

You can use the [AWS Simple Monthly Calculator](#) to estimate your monthly bill. The calculator provides per service cost breakdown, as well as an aggregate monthly estimate. You can also use the calculator to see an estimation and breakdown of costs for common solutions.

AWS has introduced a free usage tier to help you get started with AWS. Take advantage of the [Free Usage Tier](#), and get started today!

## Resources

Resource	Description
<a href="#">AWS Simple Monthly Calculator</a>	The AWS Simple Monthly Calculator helps estimate your monthly bill.
<a href="#">AWS Architecture Center</a>	The AWS Architecture Center provides you with the necessary guidance and best practices to build highly scalable and reliable applications in the AWS Cloud.
<a href="#">AWS Economics Center</a>	The AWS Economics Center provides access to information, tools, and resources to compare the costs of Amazon Web Services with IT infrastructure alternatives.
<a href="#">AWS Account Activity Page</a>	View your current charges and account activity, itemized by service and by usage type. Previous months' billing statements are also available.
<a href="#">AWS Usage Reports</a>	Usage reports are available to download for each service. Reports can be customized by specifying usage types, timeframe, service operations, and more.

## Cost Calculation Example

This section of the document uses the AWS Simple Monthly Calculator to provide costing example of a common use case of a dynamic website hosted on AWS using Amazon EC2, Auto Scaling, and Amazon RDS. The Amazon EC2 instance runs the web and application tiers, and we use Auto Scaling to match the number of instances to the traffic load. Amazon RDS uses one DB instance for its primary storage. This DB instance is deployed across multiple Availability Zones. The prices are taken as of the date of publication and may change.

### Architecture

Elastic Load Balancing balances traffic across one or more Amazon EC2 instances. The Amazon EC2 instances belong to an Auto Scaling group, and the Auto Scaling group either adds or subtracts Amazon EC2 instances, depending on variations in the traffic load. Deploying Amazon RDS across multiple Availability Zones enhances data durability and availability. Amazon RDS provisions and maintains a standby in a different Availability Zone for automatic failover in the event of planned or unplanned outages. The following illustration shows the example architecture for a dynamic website using Amazon EC2, Auto Scaling, and one Amazon RDS database instance across multiple Availability Zones.

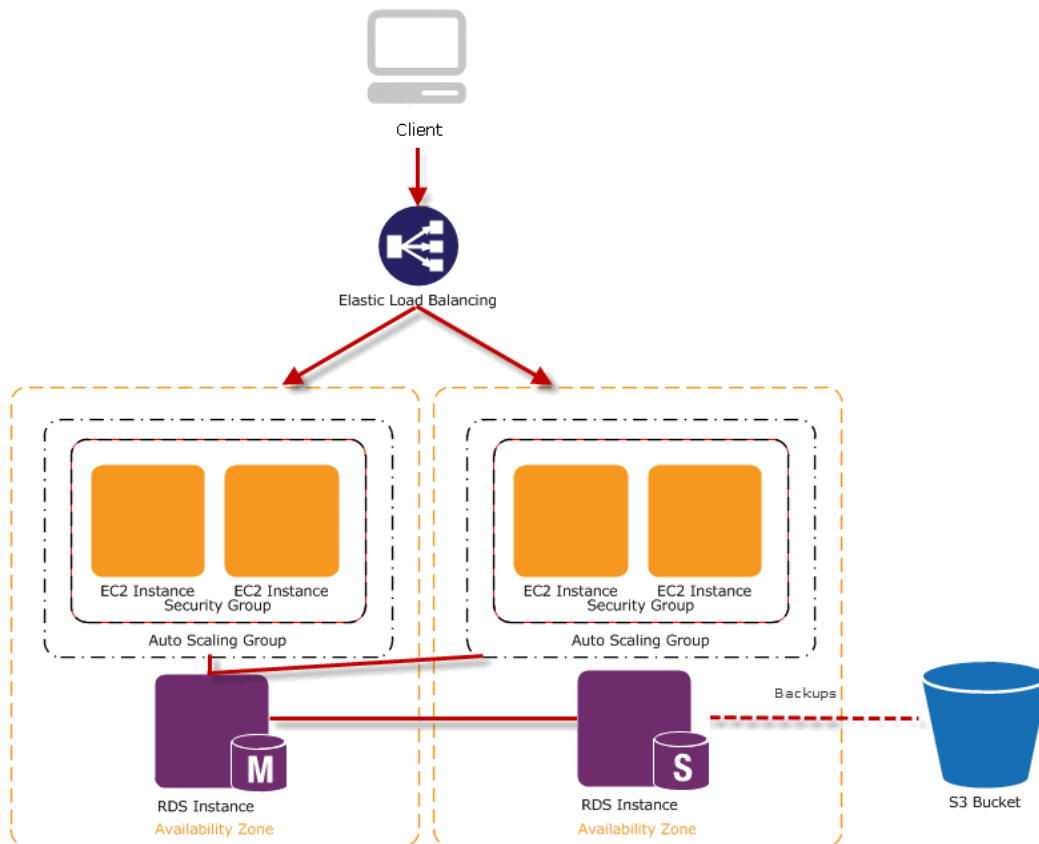


Figure 1: Complex Dynamic Site Architecture

### Daily Usage Profile

You can monitor daily usage for your application so that you can better estimate your costs. For instance, you can look at the daily pattern to figure out how your application handles traffic. For each hour, track how many hits you get at

your website. Then track how many instances are running. Add up the total number of hits for that day. Examine the number of Amazon EC2 instances that run each hour, and then take the average. You can use the number of hits per day and the average number of instances for your calculations.

### Amazon EC2 Cost Breakdown

The following table shows the characteristics for Amazon EC2 we have identified for this dynamic site.

Characteristic	Estimated Usage	Description
Clock Hours of Server Time	24 hrs/day	At an average of 30.5 days in a month, the instance runs 732 hours/month
Machine Characteristics	m3.medium instance, instance storage	3.75 GiB of memory, 1 vCPU, 4 GB of SSD storage
Additional Storage	No EBS volumes	For this application, we need only the local storage which is packaged with the AMI.
Data Transfer	Data In: 5 GB/day Data Out: 50 GB /day	There are approximately 1,000,000 hits per day. Each response transfers out about 50 KB, and each request transfers in about 5 KB.
Instance Scale	4	On average in a given day, there are 4 instances running.
Elastic Load Balancing	Hourly usage: 732 hrs/month Data processed: 1677.5 GB/month	Elastic Load Balancing is used 24 hours/day, 7 days/week. Elastic Load Balancing processes a total of 55 GB/day (data in + data out)
Elastic IP Address	None	We use Elastic Load Balancing to balance traffic across instances.
Detailed Monitoring	None	Detailed monitoring is not enabled, so we get basic monitoring free of charge.

The total cost for one month is the sum of the cost of the running instances, overall AWS data transfer, Elastic Load Balancers, and the data processed by the Elastic Load Balancers.

Variable	Formula	Calculation
<b>Amazon EC2 Instance Cost</b>	$\frac{\text{Instance cost per hour} \times \text{Number of instances}}{\text{Uptime in hours}}$	$\frac{\$0.070 \times 4}{732} = \$204.96$
<b>Elastic Load Balancer Cost</b>	$\frac{(\text{Hours used} \times \text{Hourly Rate}) + (\text{Data processed (GB)} \times \text{Process Rate})}{1}$	$\frac{732 \times \$0.025 + 1677.5 \times \$0.008}{1} = \$31.72$
<b>AWS Data Transfer Cost</b>	$\frac{(\text{Data in (GB)} \times \text{Data In Rate}) + (\text{Data out (GB)} \times \text{Data Out Rate})}{1}$	$\frac{152.5 \times \$0.00 + (1525-1)^1 \times \$0.12}{1} = \$182.88$
<b>Estimated Cost for Amazon EC2</b>		<b>\$419.56</b>

Prices shown for US East Region as of July 20<sup>th</sup> 2014

<sup>1</sup> Since up to one gigabyte is free per month, we subtract one.

We use the [AWS Simple Monthly Calculator](#) to estimate this cost. Using the calculator, we arrive at the total cost for one month's usage for Amazon EC2 of \$419.56<sup>2</sup>.

### Amazon RDS Cost Breakdown

The following table shows the characteristics for Amazon RDS we have identified for this complex dynamic site.

Characteristic	Estimated Usage	Description
<b>Clock Hours of Server Time</b>	24 hrs/day	24*30.5 = 732 hours per month
<b>Database Characteristics</b>	Medium RDS instance	3.75 GiB memory, 1 virtual core, Moderate network performance
<b>Provisioned Storage</b>	100 GB/month	Amazon provides 5 GB to 1 TB of associated storage capacity for your primary data set.
<b>Requests</b>	300M I/O requests/month	We have 1,000,000 hits per day at a rate of 5 I/O requests per hit on site with 30.5 days in a month on average. This is a total of 150M I/O requests per month, but since the write I/O request will double because data is also replicated to the standby instance, we have a total of 300M.
<b>Deployment Type</b>	Multiple Availability Zones	We will run our database instance across multiple Availability Zones.
<b>Additional Backup Storage</b>	None	We'll use up to the provisioned amount, which is 100 GB.
<b>Data Transfer</b>	Data in: 0 GB Data out: 0 GB	There is no data transfer from RDS to the Internet.
<b>Database Instance Scale</b>	1	We need one database instance.
<b>Special Features</b>	None	There are no reserved database instances.

Because we do not have data transfer in or out, and we do not have additional backup storage, the total cost for one month is the sum of the cost of the running instances, provisioned storage, and I/O requests.

Variable	Formula	Calculation
<b>Instance Cost</b>	$\text{Instance cost per hour} \times \text{Number of instances} \times \text{Uptime in hours}$	$\$0.180 \times 1 \times 732 = \$131.76$
<b>Provisioned Storage</b>	$\text{Storage rate} \times \text{Storage Amount (GB)}$	$\$0.20 \times 100 = \$20.00$
<b>I/O Requests</b>	$\text{I/O rate} \times \text{Number of requests}$	$\$0.10 \times 300 = \$30.00$

<sup>2</sup> Depending on regional factors and special offers, the cost you get from the calculator will be slightly different.

Estimated cost for Amazon RDS		\$181.76
-------------------------------	--	----------

*Prices shown for US East Region as of July 20<sup>th</sup> 2014*

We use the [AWS Simple Monthly Calculator](#) to estimate this cost. The total cost for one month's usage for Amazon RDS portion of this scenario is \$181.76<sup>3</sup>.

### Total Cost

To calculate the total cost for this example, we add the cost for Amazon EC2, Amazon RDS, and AWS Data Transfer Out and subtract any discount that falls into the AWS Free Usage Tier. The total cost of the dynamic site in this scenario is estimated at \$601.32<sup>3</sup> per month.

---

<sup>3</sup> Depending on regional factors and special offers, the cost you get from the [AWS Simple Monthly Calculator](#) will be slightly different.